
Introduction to Search




Fundamentals of Artificial Intelligence

Slides are mostly adapted from
AIMA, Svetlana Lazebnik (UIUC) and Percy Liang (Stanford)

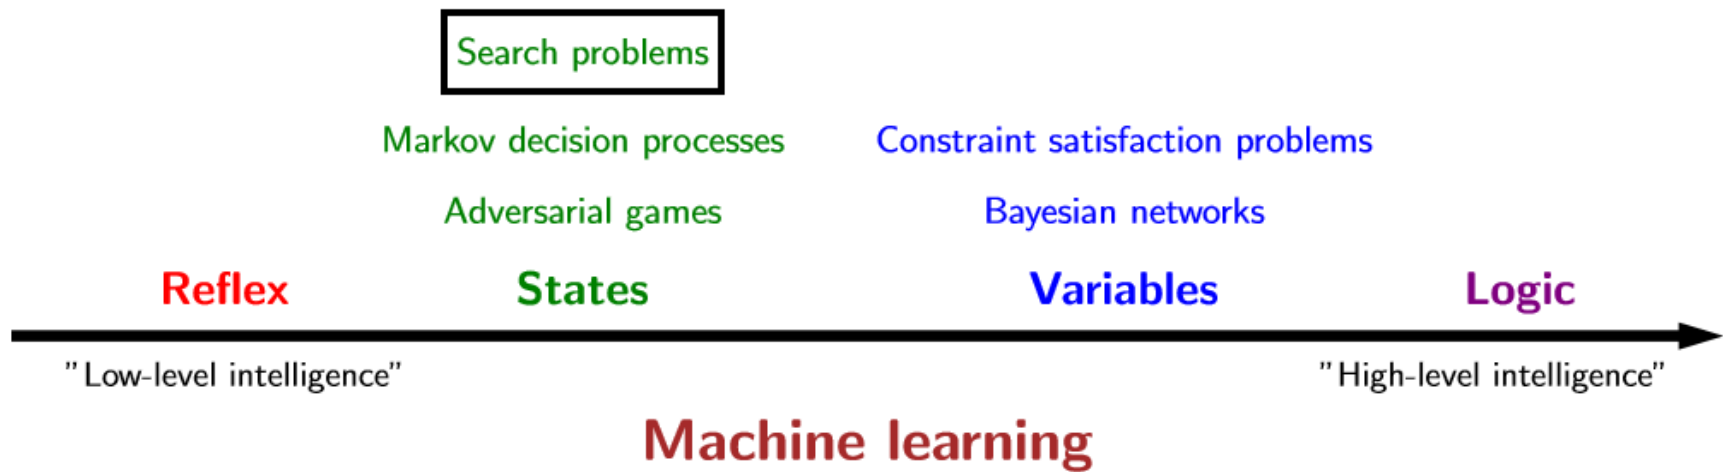
A farmer wants to get his cabbage, goat, and wolf across a river. He has a boat that only holds two. He cannot leave the cabbage and goat alone or the goat and wolf alone. How many river crossings does he need

When you solve this problem, try to think about how you did it. You probably simulated the scenario in your head, trying to send the farmer over with the goat, observing the consequences. If nothing got eaten, you might continue with the next action. Otherwise, you undo that move and try something else.


How can we get a machine to do this automatically? One of the things we need is a systematic approach that considers all the possibilities. We will see that search problems define the possibilities, and search algorithms explore these possibilities.



Sometimes you can do better if you change the model (perhaps the value of having a wolf is zero) instead of focusing on the algorithm.



Application: route finding



Objective: shortest? fastest? most scenic?

Actions: go straight, turn left, turn right


Application: robot motion planning



Objective: fastest? most energy efficient? safest?

Actions: translate and rotate joints

Application: solving puzzles



Objective: reach a certain configuration

Actions: move pieces (e.g., Move12Down)

Types of agents

Reflex agent



- Consider how the world **IS**
- Choose action based on current percept
- Do not consider the future consequences of actions



Planning agent



- Consider how the world **WOULD BE**
- Decisions based on (hypothesized) consequences of actions
- Must have a model of how the world evolves in response to actions
- Must formulate a goal



Reflex Agents

- Reflex Agents:
 - Choose action based on current percept (and maybe memory)
 - May have memory or a model of the world's current state
 - Do not consider the future consequences of their actions
 - Act on how the world IS
- Can a reflex agent be rational?



Goal Based Agents

- Goal-based agents:
 - Plan ahead
 - Decisions based on (hypothesized) consequences of actions
 - Must have a model of how the world evolves in response to actions
 - Act on how the world **WOULD BE**



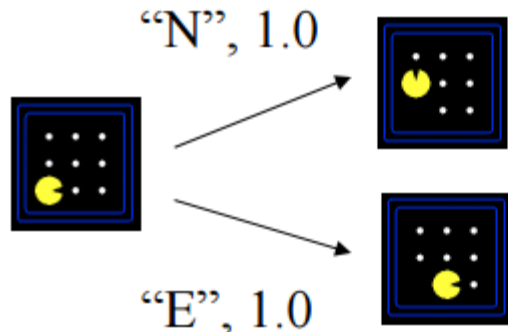
Search Problems

- A search problem consists of:

- A state space



- A successor function




- A start state and a goal test

- A solution is a sequence of actions (a plan) which transforms the start state to a goal state
 - The **performance measure** is defined by (a) reaching the goal and (b) how “expensive” the path to the goal is

Search

- We will consider the problem of designing **goal-based agents** in **fully observable, deterministic, discrete, known** environments

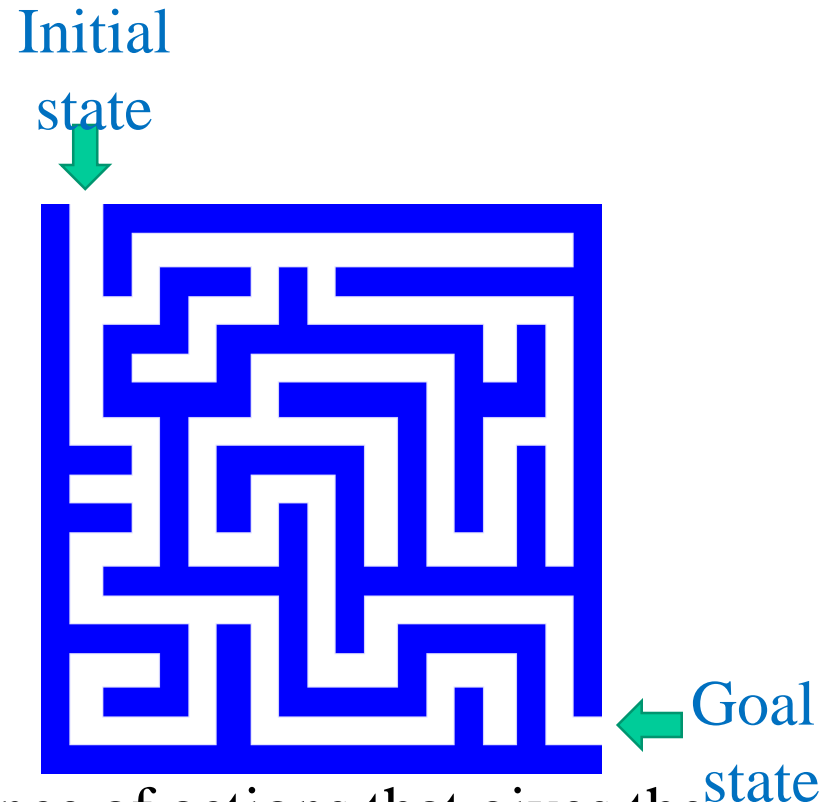


Search

- We will consider the problem of designing **goal-based agents** in **fully observable, deterministic, discrete, known** environments
 - The agent must find a *sequence of actions* that reaches the goal
 - The **performance measure** is defined by (a) reaching the goal and (b) how “expensive” the path to the goal is
 - We are focused on the process of finding the solution; while executing the solution, we assume that the agent can safely ignore its percepts (**open-loop system**)

Search problem components

- **Initial state**
- **Actions**
- **Transition model**
 - What state results from performing a given action in a given state?
- **Goal state**
- **Path cost**
 - Assume that it is a sum of nonnegative *step costs*
- The **optimal solution** is the sequence of actions that gives the *lowest* path cost for reaching the goal



Example: Romania

- On vacation in Romania; currently in Arad
- Flight leaves tomorrow from Bucharest

- **Initial state**

- Arad

- **Actions**

- Go from one city to another

- **Transition model**


- If you go from city A to city B, you end up in city B

- **Goal state**

- Bucharest

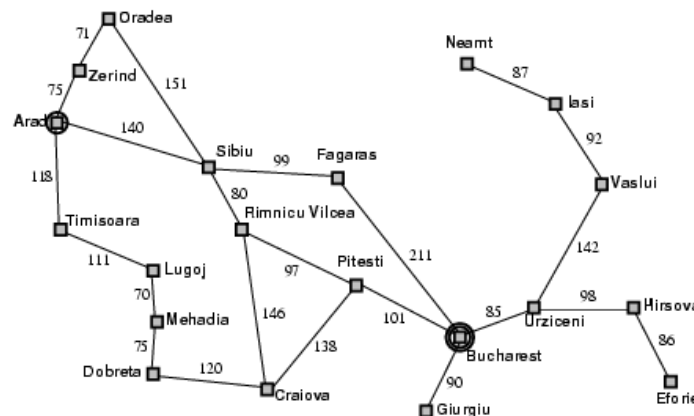
- **Path cost**

- Sum of edge costs (total distance traveled)




State space

- The initial state, actions, and transition model define the **state space** of the problem
 - The set of all states reachable from initial state by any sequence of actions
 - Can be represented as a **directed graph** where the nodes are states and links between nodes are actions
- What is the state space for the Romania problem?




Example: Vacuum world



- **States**
 - Agent location and dirt location
 - How many possible states?
 - What if there are n possible locations?
 - The size of the state space grows exponentially with the “size” of the world!
- **Actions**
 - Left, right, suck
- **Transition model**

Vacuum world state space graph



Example: The 8-puzzle

- **States**

- Locations of tiles

- 8-puzzle: 181,440 states ($9!/2$)
- 15-puzzle: ~10 trillion states
- 24-puzzle: $\sim 10^{25}$ states

- **Actions**

- Move blank left, right, up, down

- **Path cost**

- 1 per move

7	2	4
5		6
8	3	1


Start State

	1	2
3	4	5
6	7	8

Goal State

- Finding the optimal solution of n-Puzzle is NP-hard

Example: Robot motion planning



- **States**
 - Real-valued joint parameters (angles, displacements)
- **Actions**
 - Continuous motions of robot joints
- **Goal state**
 - Configuration in which object is grasped
- **Path cost**
 - Time to execute, smoothness of path, etc.

Search

- Given:
 - **Initial state**
 - **Actions**
 - **Transition model**
 - **Goal state**
 - **Path cost**
- How do we find the optimal solution?
 - How about building the state space and then using Dijkstra's shortest path algorithm?
 - Complexity of Dijkstra's is $O(E + V \log V)$, where V is the size of the state space
 - The state space may be huge!


Search: Basic idea

- Let's begin at the start state and **expand** it by making a list of all possible successor states
- Maintain a **frontier** or a list of unexpanded states
- At each step, pick a state from the frontier to expand
- Keep going until you reach a goal state
- Try to expand as few states as possible


Search: Basic idea




Search: Basic idea




Search: Basic idea




Search: Basic idea




Search: Basic idea




Search: Basic idea




Search: Basic idea




Search: Basic idea




Search: Basic idea




Search: Basic idea



Search: Basic idea



Search: Basic idea



Tree Search Algorithm Outline

- Initialize the **frontier** using the **starting state**
- While the frontier is not empty
 - Choose a frontier node according to **search strategy** and take it off the frontier
 - If the node contains the **goal state**, return solution
 - Else **expand** the node and add its children to the frontier




Farmer Cabbage Goat Wolf

Actions:

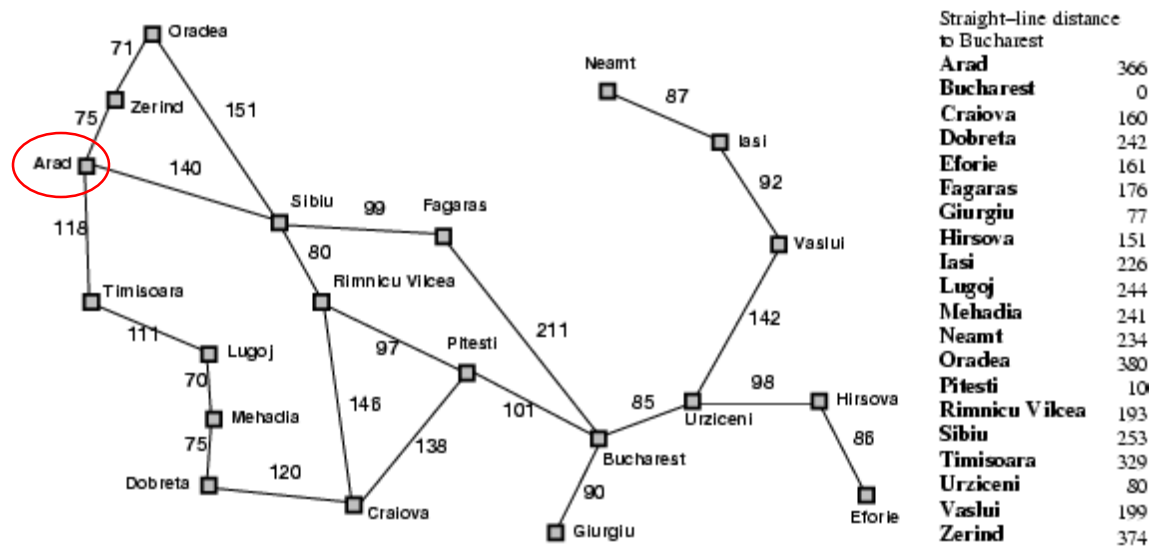
F▷	F◁
FC▷	FC◁
FG▷	FG◁
FW▷	FW◁

Approach: build a **search tree** ("what if?")




Tree search example

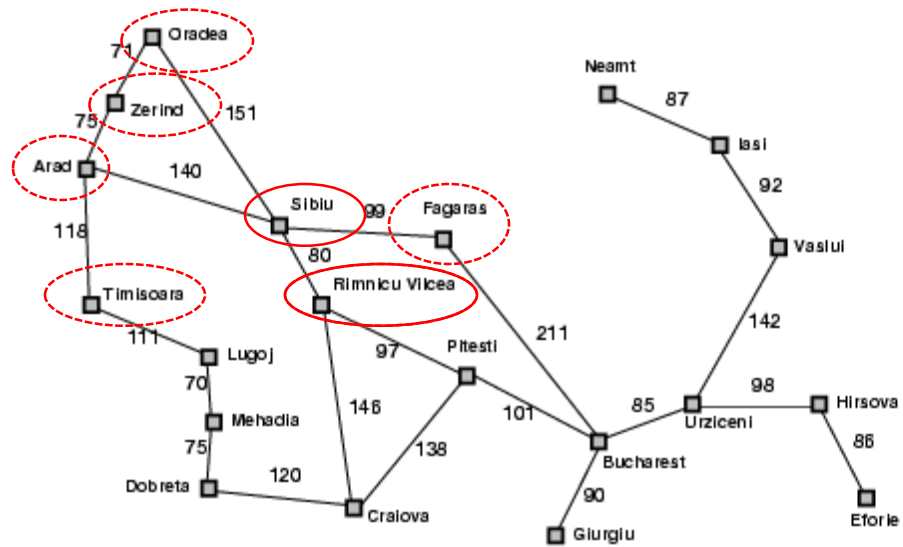

Arad
366=0+366



Tree search example



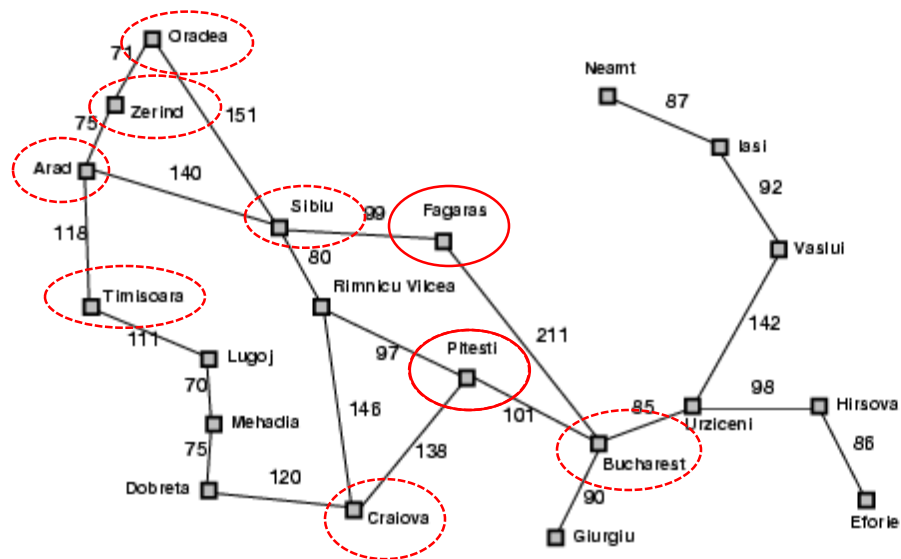

Tree search example



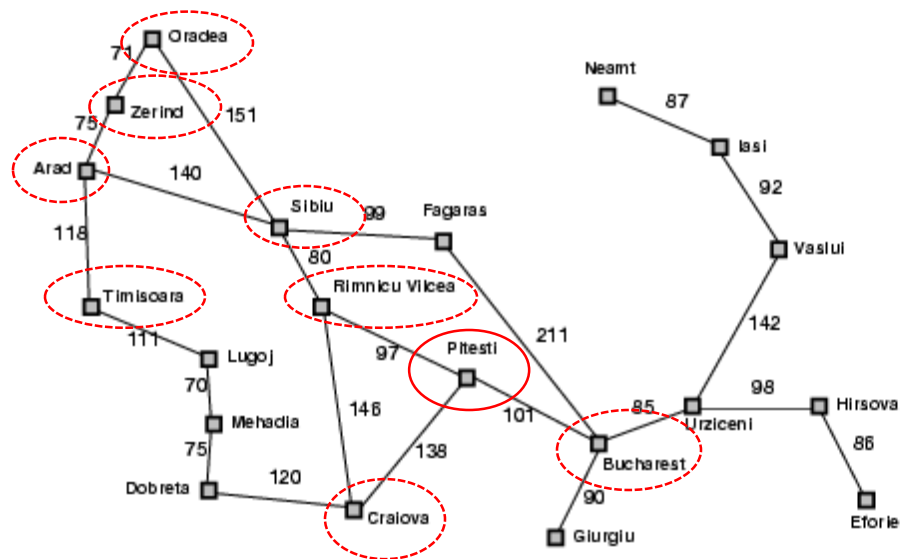

Straight-line distance
to Bucharest

Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	176
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	10
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

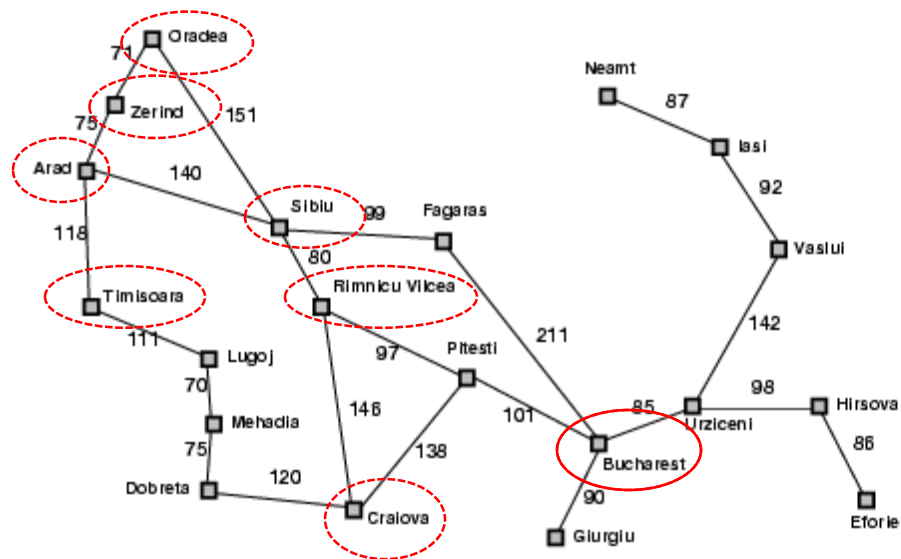

Tree search example



Tree search example



Tree search example



Straight-line distance
to Bucharest

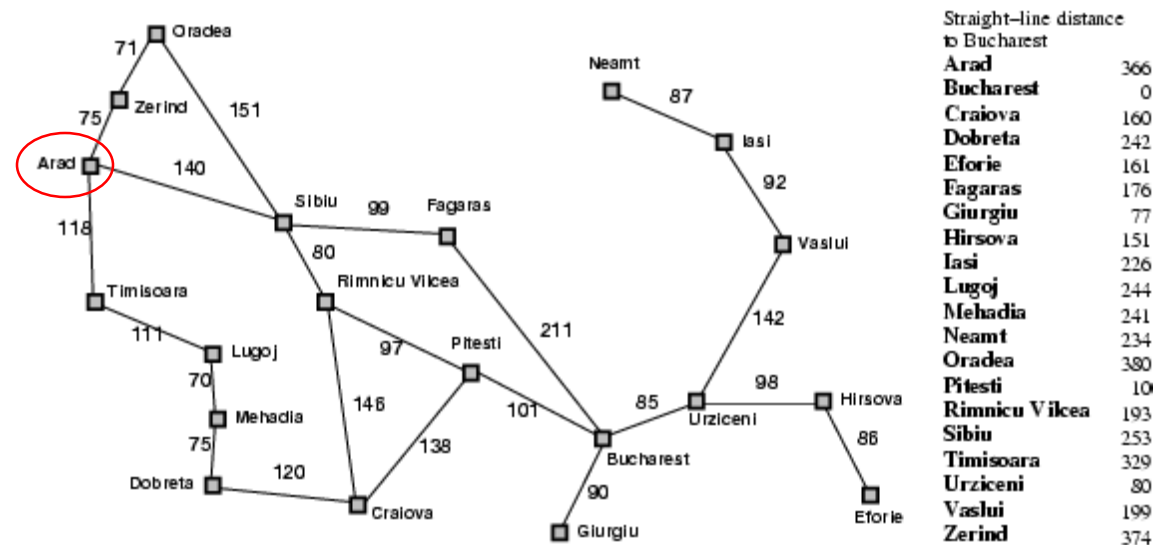
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	176
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	10
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

Handling repeated states


- Initialize the **frontier** using the **starting state**
- While the frontier is not empty
 - Choose a frontier node according to **search strategy** and take it off the frontier
 - If the node contains the **goal state**, return solution
 - Else **expand** the node and add its children to the frontier
- To handle repeated states:
 - Every time you expand a node, add that state to the **explored set**; do not put explored states on the frontier again
 - Every time you add a node to the frontier, check whether it already exists in the frontier with a higher path cost, and if yes, replace that node with the new one

Search without repeated states

Arad
366=0+366



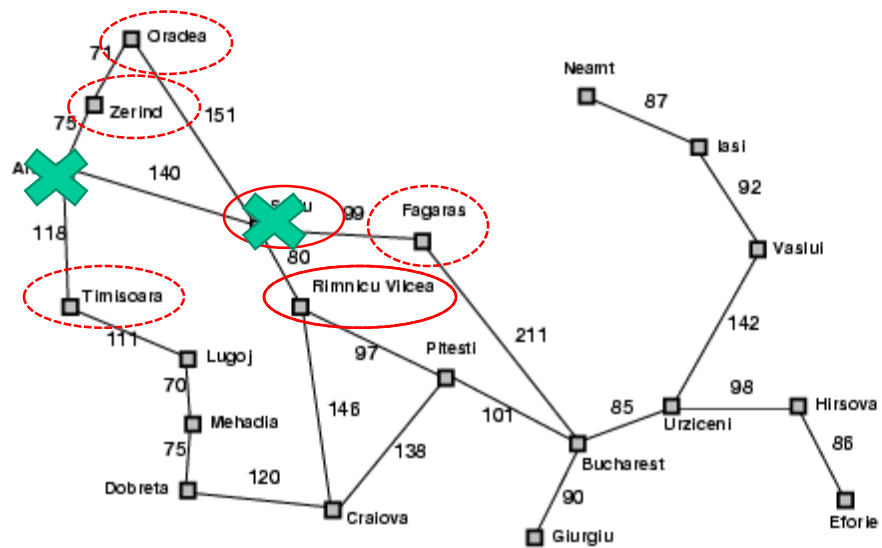

Search without repeated states



Straight-line distance to Bucharest

Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	176
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	10
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

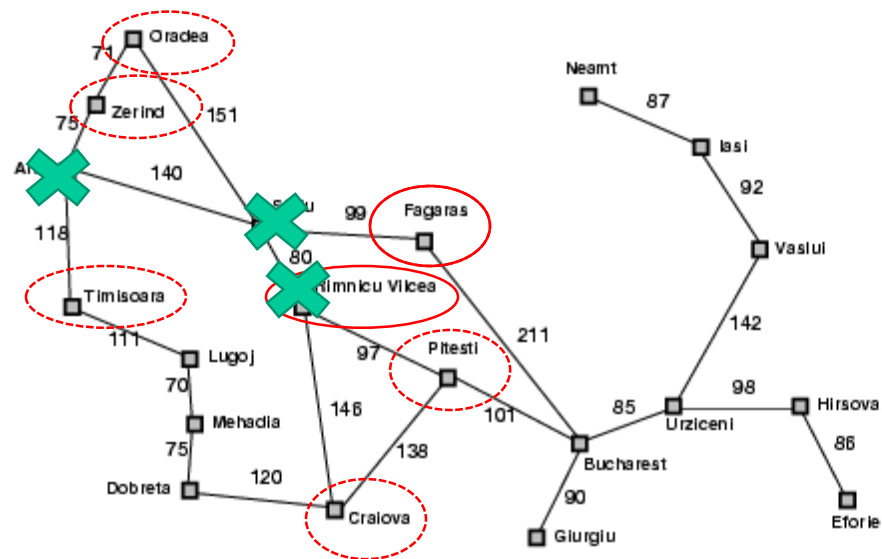
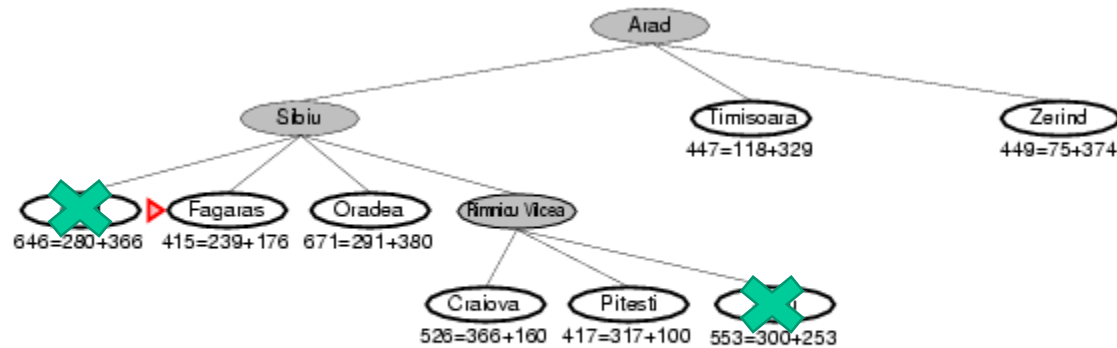
Search without repeated states



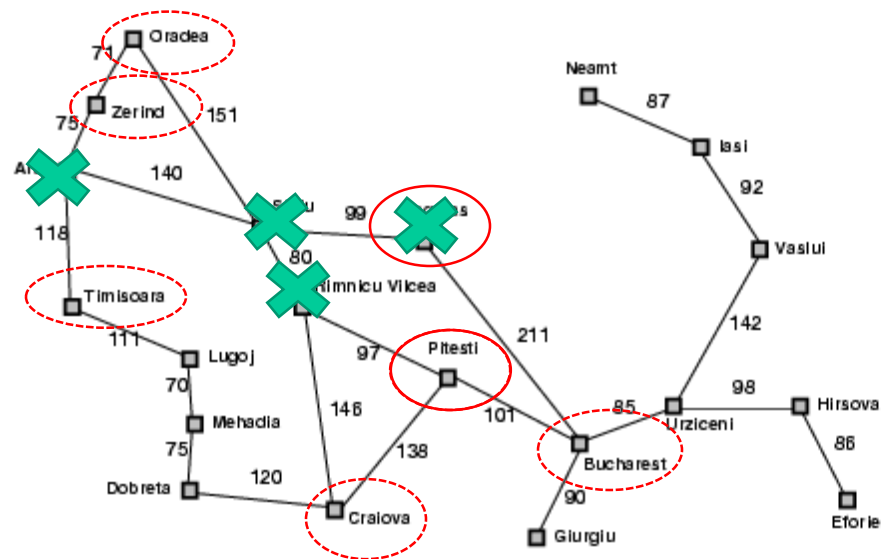
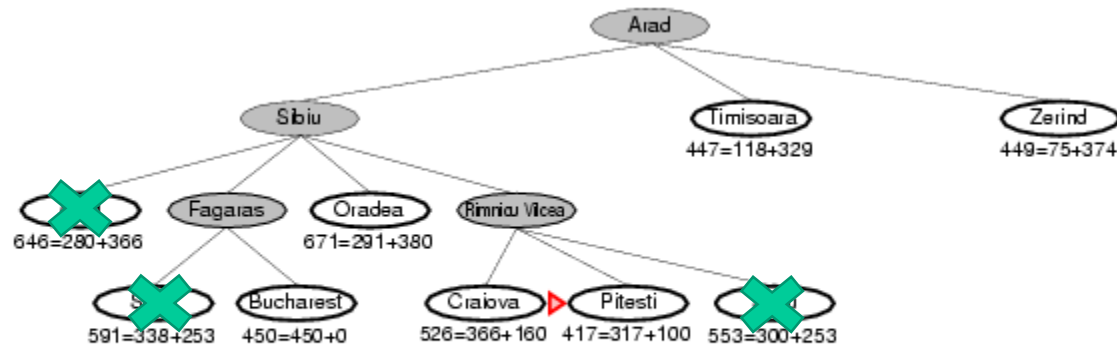
Straight-line distance to Bucharest

Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	176
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	10
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

Search without repeated states



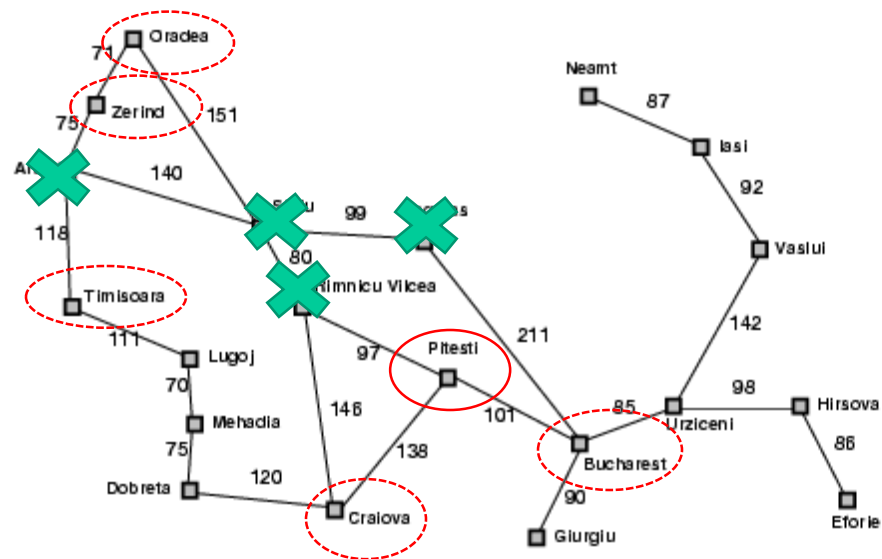
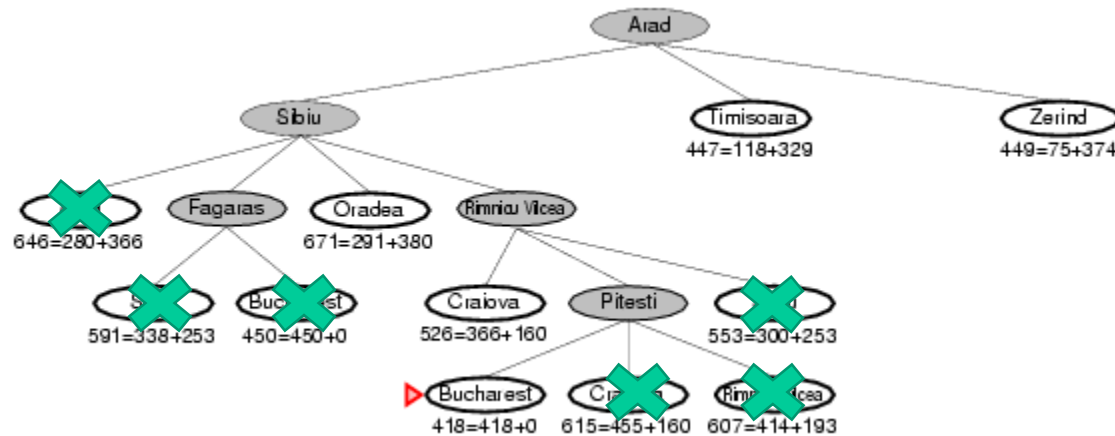
Search without repeated states



Straight-line distance
to Bucharest

Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	176
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	10
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

Search without repeated states



Search without repeated states

