

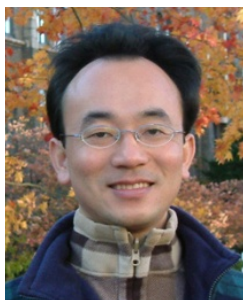
人工智慧 (Artificial Intelligence)

問題解決 (Problem Solving)

1092AI03

MBA, IM, NTPU (M5010) (Spring 2021)

Wed 2, 3, 4 (9:10-12:00) (B8F40)



Min-Yuh Day

戴敏育

Associate Professor

副教授

Institute of Information Management, National Taipei University

國立臺北大學 資訊管理研究所

<https://web.ntpu.edu.tw/~myday>

2021-03-09



課程大綱 (Syllabus)

- | 週次 (Week) | 日期 (Date) | 內容 (Subject/Topics) |
|-----------|------------|--|
| 1 | 2021/02/24 | 人工智慧概論 (Introduction to Artificial Intelligence) |
| 2 | 2021/03/03 | 人工智慧和智慧代理人 (Artificial Intelligence and Intelligent Agents) |
| 3 | 2021/03/10 | 問題解決 (Problem Solving) |
| 4 | 2021/03/17 | 知識推理和知識表達 (Knowledge, Reasoning and Knowledge Representation) |
| 5 | 2021/03/24 | 不確定知識和推理 (Uncertain Knowledge and Reasoning) |
| 6 | 2021/03/31 | 人工智慧個案研究 I (Case Study on Artificial Intelligence I) |

課程大綱 (Syllabus)

| 週次 (Week) | 日期 (Date) | 內容 (Subject/Topics) |
|-----------|------------|---|
| 7 | 2021/04/07 | 放假一天 (Day off) |
| 8 | 2021/04/14 | 機器學習與監督式學習 (Machine Learning and Supervised Learning) |
| 9 | 2021/04/21 | 期中報告 (Midterm Project Report) |
| 10 | 2021/04/28 | 學習理論與綜合學習 (The Theory of Learning and Ensemble Learning) |
| 11 | 2021/05/05 | 深度學習 (Deep Learning) |
| 12 | 2021/05/12 | 人工智慧個案研究 II (Case Study on Artificial Intelligence II) |

課程大綱 (Syllabus)

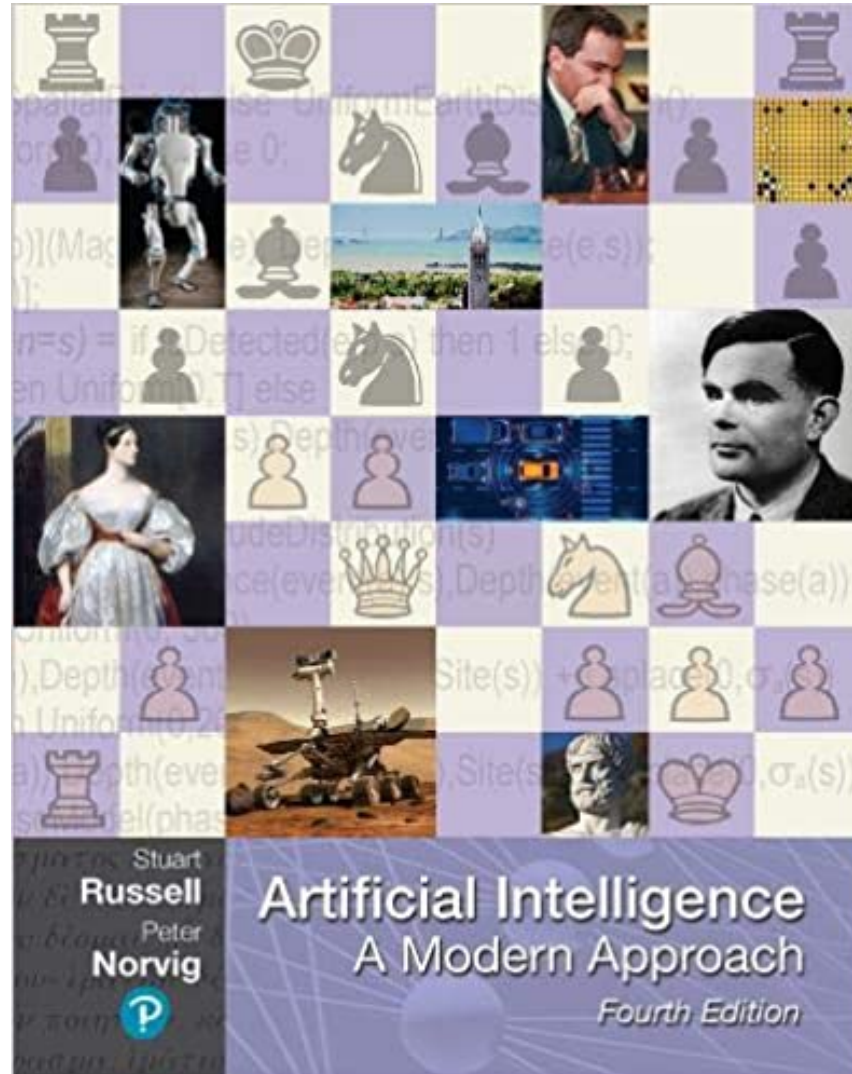
- | 週次 (Week) | 日期 (Date) | 內容 (Subject/Topics) |
|-----------|------------|--|
| 13 | 2021/05/19 | 強化學習 (Reinforcement Learning) |
| 14 | 2021/05/26 | 深度學習自然語言處理 (Deep Learning for Natural Language Processing) |
| 15 | 2021/06/02 | 機器人技術 (Robotics) |
| 16 | 2021/06/09 | 人工智慧哲學與倫理，人工智慧的未來 (Philosophy and Ethics of AI, The Future of AI) |
| 17 | 2021/06/16 | 期末報告 I (Final Project Report I) |
| 18 | 2021/06/23 | 期末報告 II (Final Project Report II) |

Artificial Intelligence Problem Solving

Outline

- **Solving Problems by Searching**
- **Search in Complex Environments**
- **Adversarial Search and Games**
- **Constraint Satisfaction Problems**

Stuart Russell and Peter Norvig (2020),
Artificial Intelligence: A Modern Approach,
4th Edition, Pearson



Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

<https://www.amazon.com/Artificial-Intelligence-A-Modern-Approach/dp/0134610997/>

Artificial Intelligence: A Modern Approach

1. Artificial Intelligence
2. Problem Solving
3. Knowledge and Reasoning
4. Uncertain Knowledge and Reasoning
5. Machine Learning
6. Communicating, Perceiving, and Acting
7. Philosophy and Ethics of AI

Artificial Intelligence: Problem Solving

Artificial Intelligence:

2. Problem Solving

- Solving Problems by Searching
- Search in Complex Environments
- Adversarial Search and Games
- Constraint Satisfaction Problems

Intelligent Agents

4 Approaches of AI

2.

**Thinking Humanly:
The Cognitive
Modeling Approach**

3.

**Thinking Rationally:
The “Laws of Thought”
Approach**

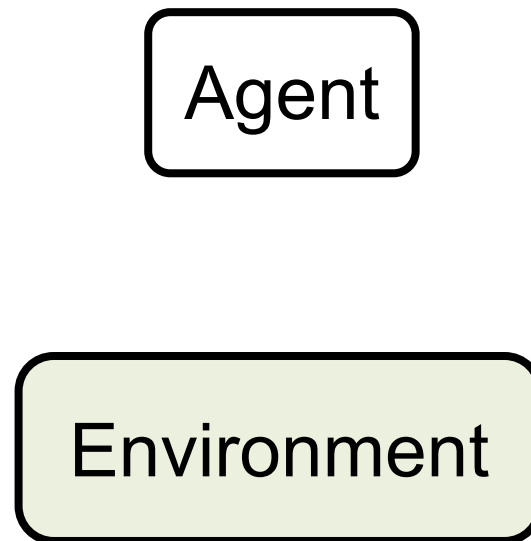
1.

**Acting Humanly:
The Turing Test
Approach** (1950)

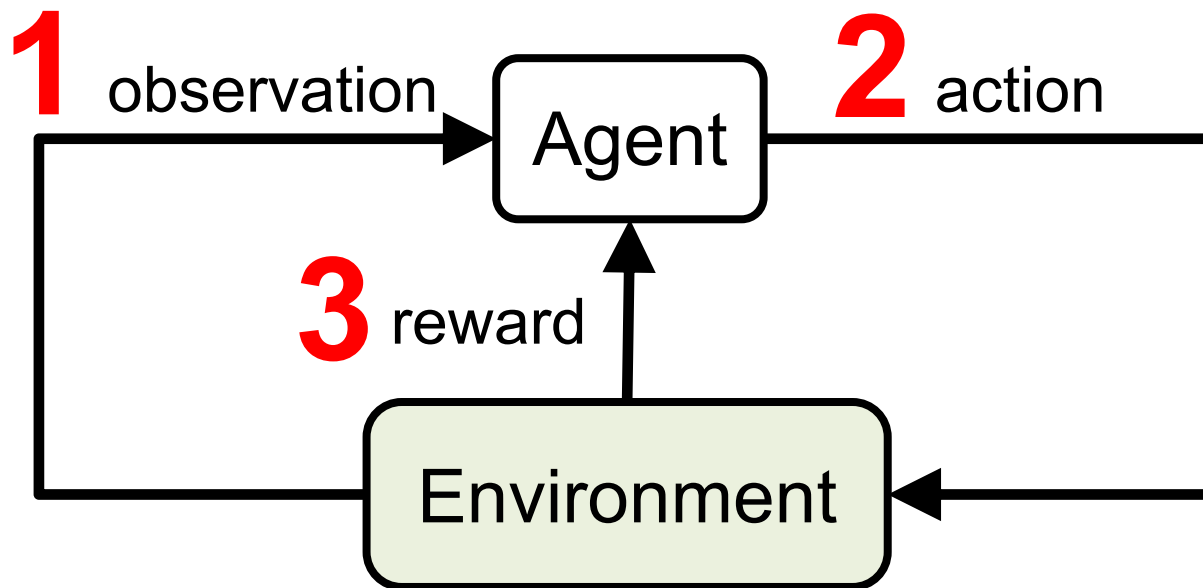
4.

**Acting Rationally:
The Rational Agent
Approach**

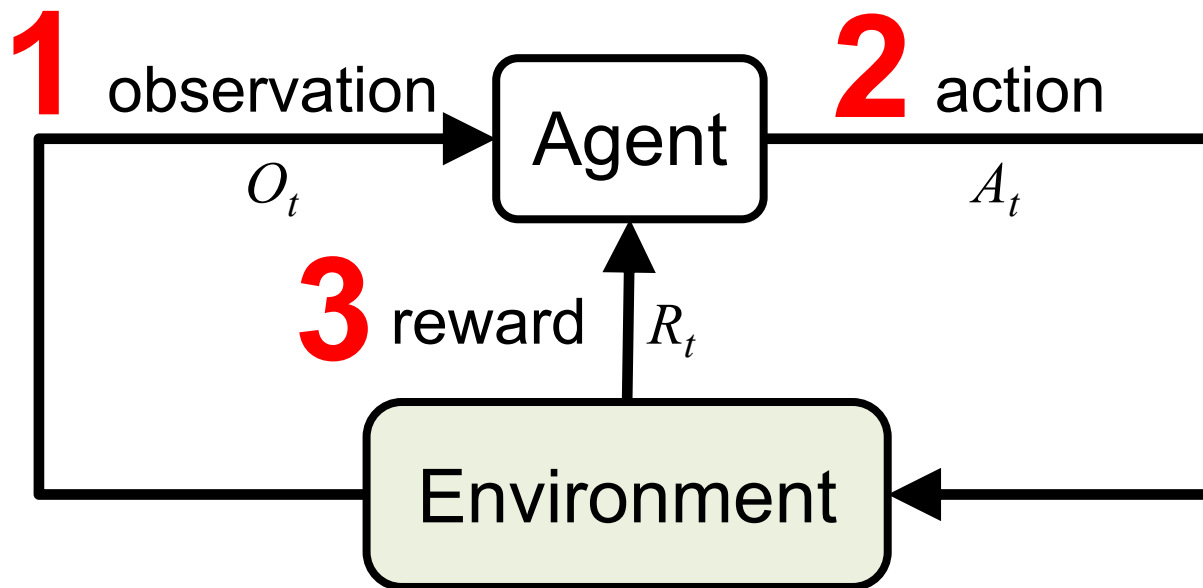
Reinforcement Learning (DL)



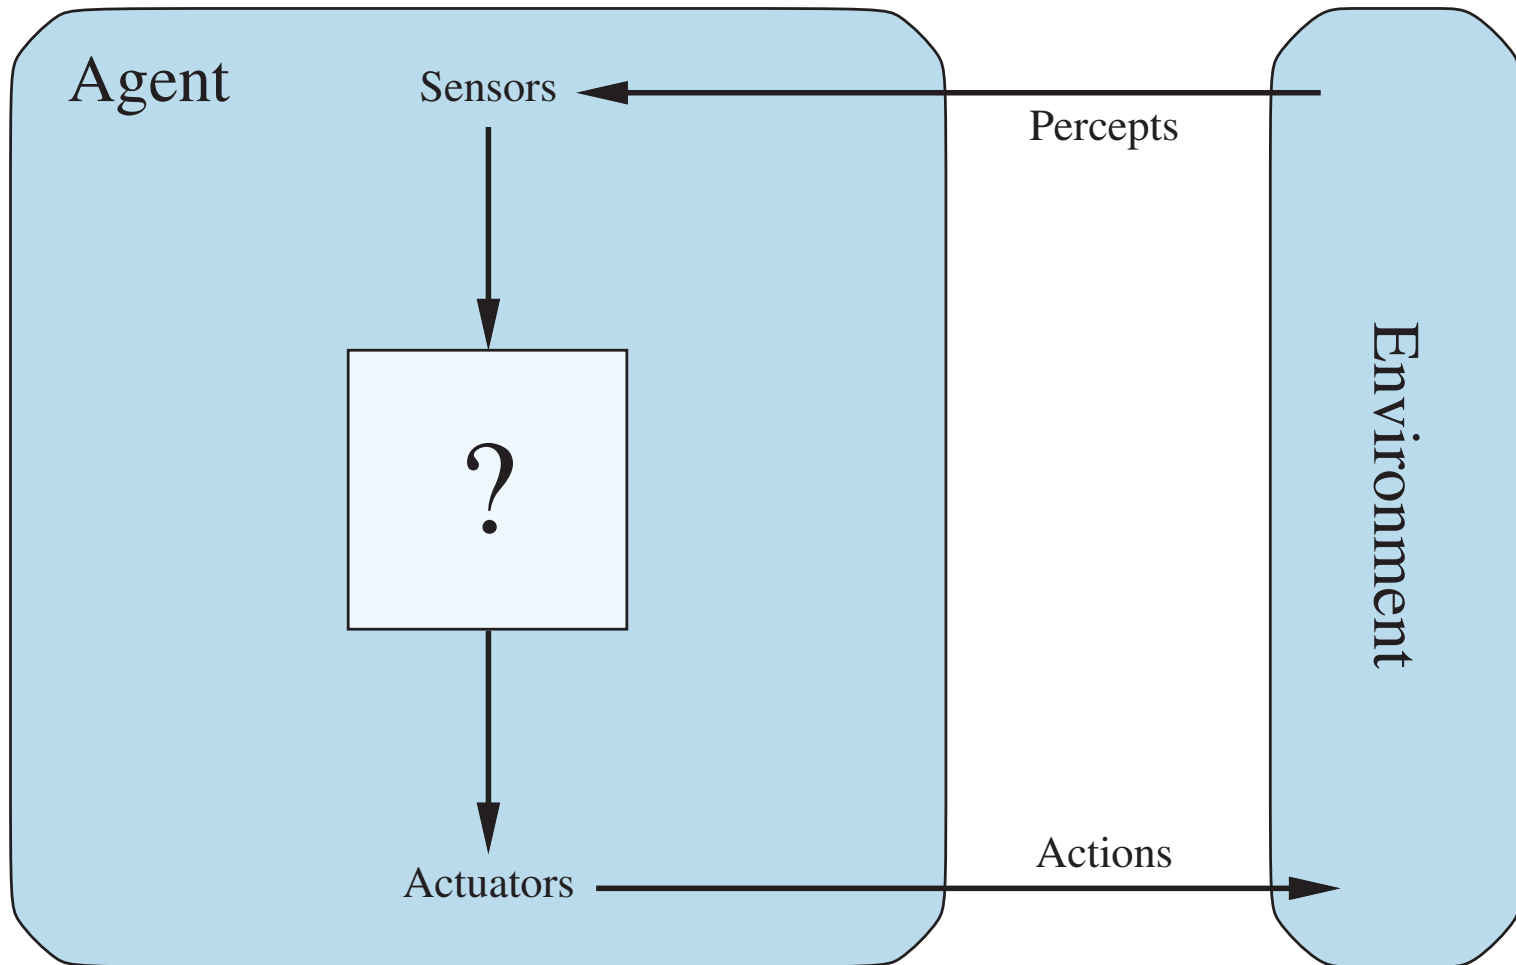
Reinforcement Learning (DL)



Reinforcement Learning (DL)



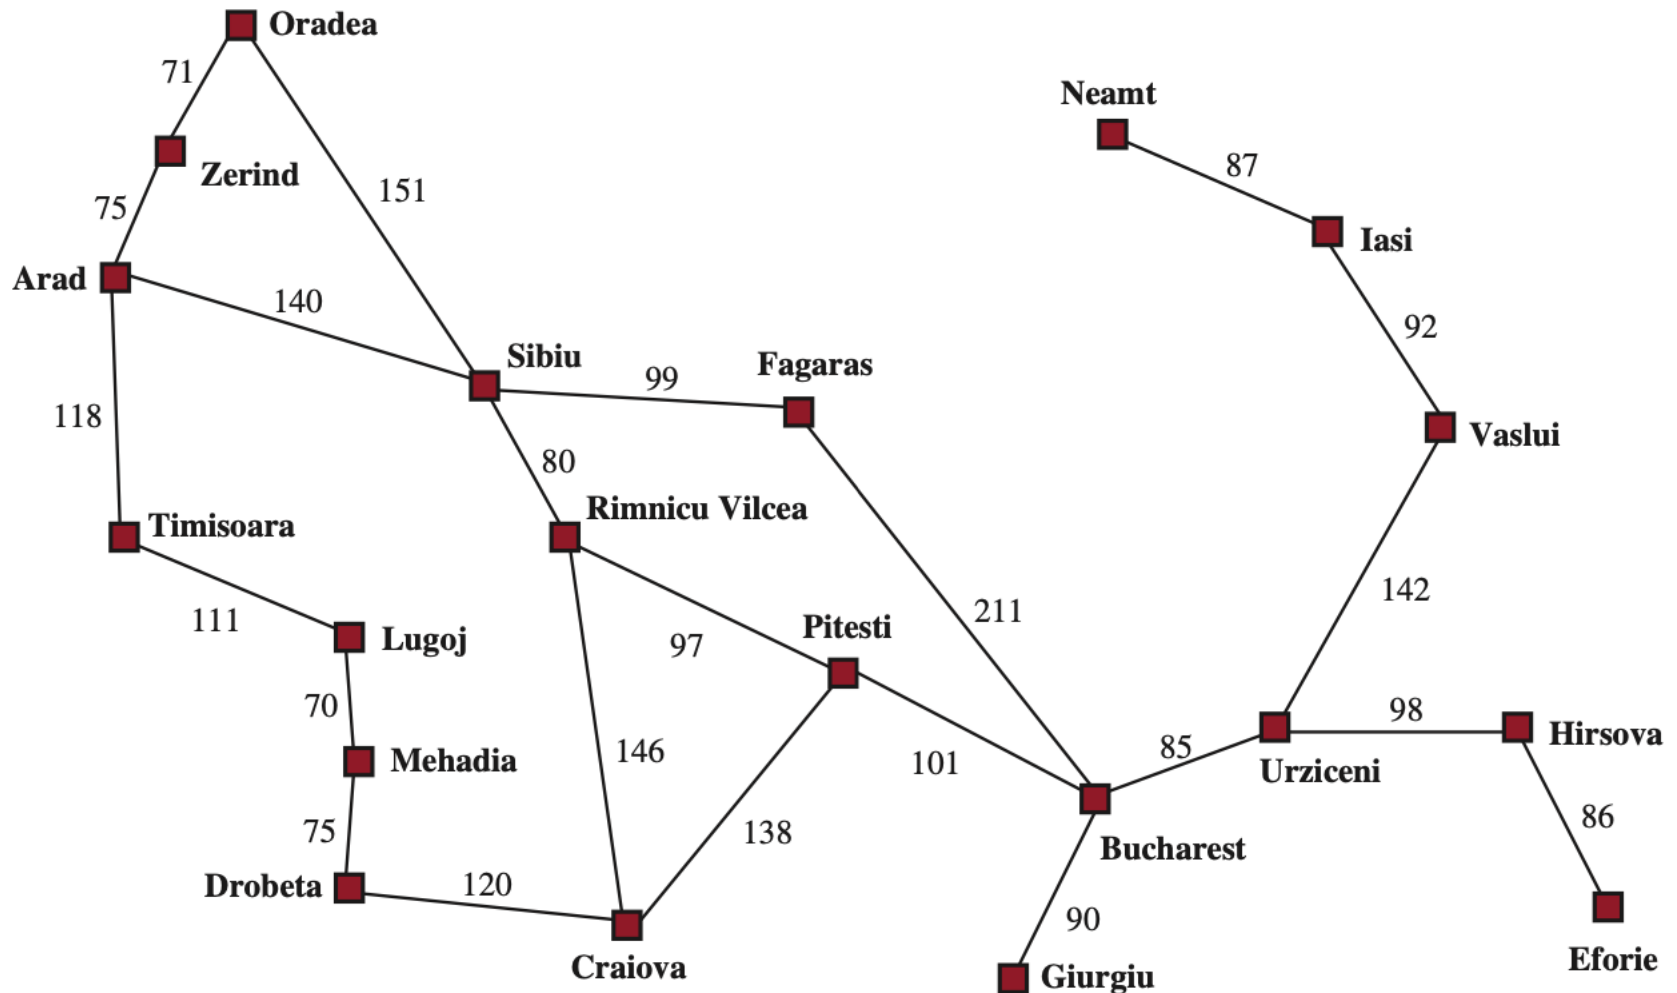
Agents interact with environments through sensors and actuators



Solving Problems by Searching

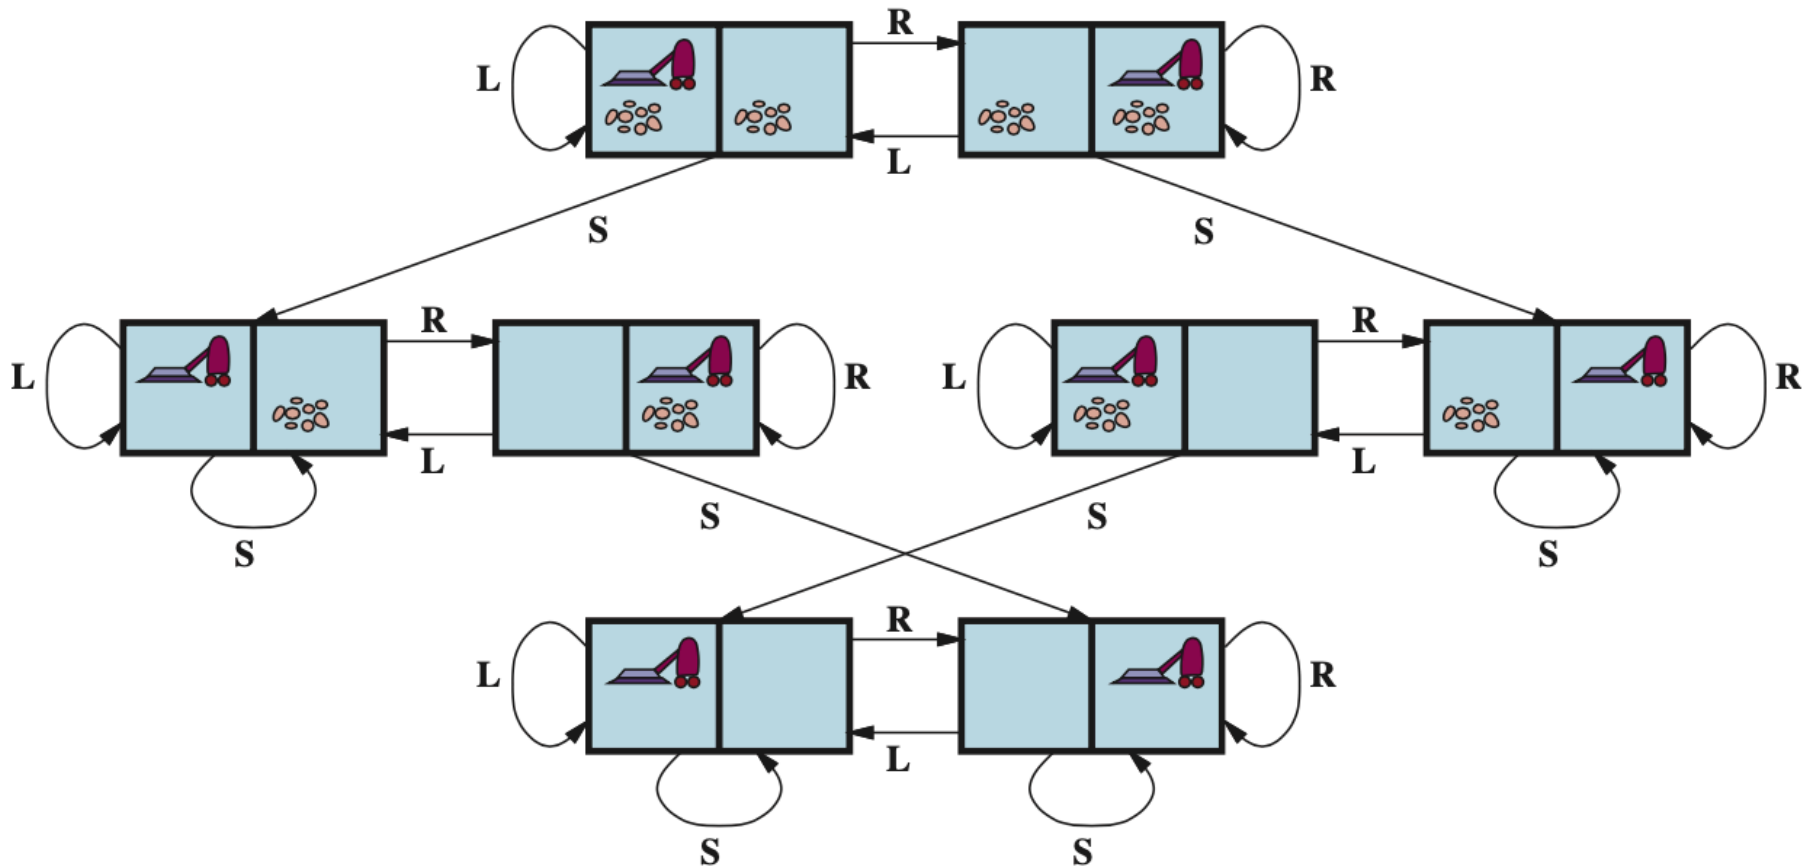
AI: Solving Problems by Searching

A simplified road map of part of Romania, with road distances in miles.



The state-space graph for the two-cell vacuum world

There are 8 states and three actions for each state:
L = Left, R = Right, S = Suck.



A typical instance of the 8-puzzle

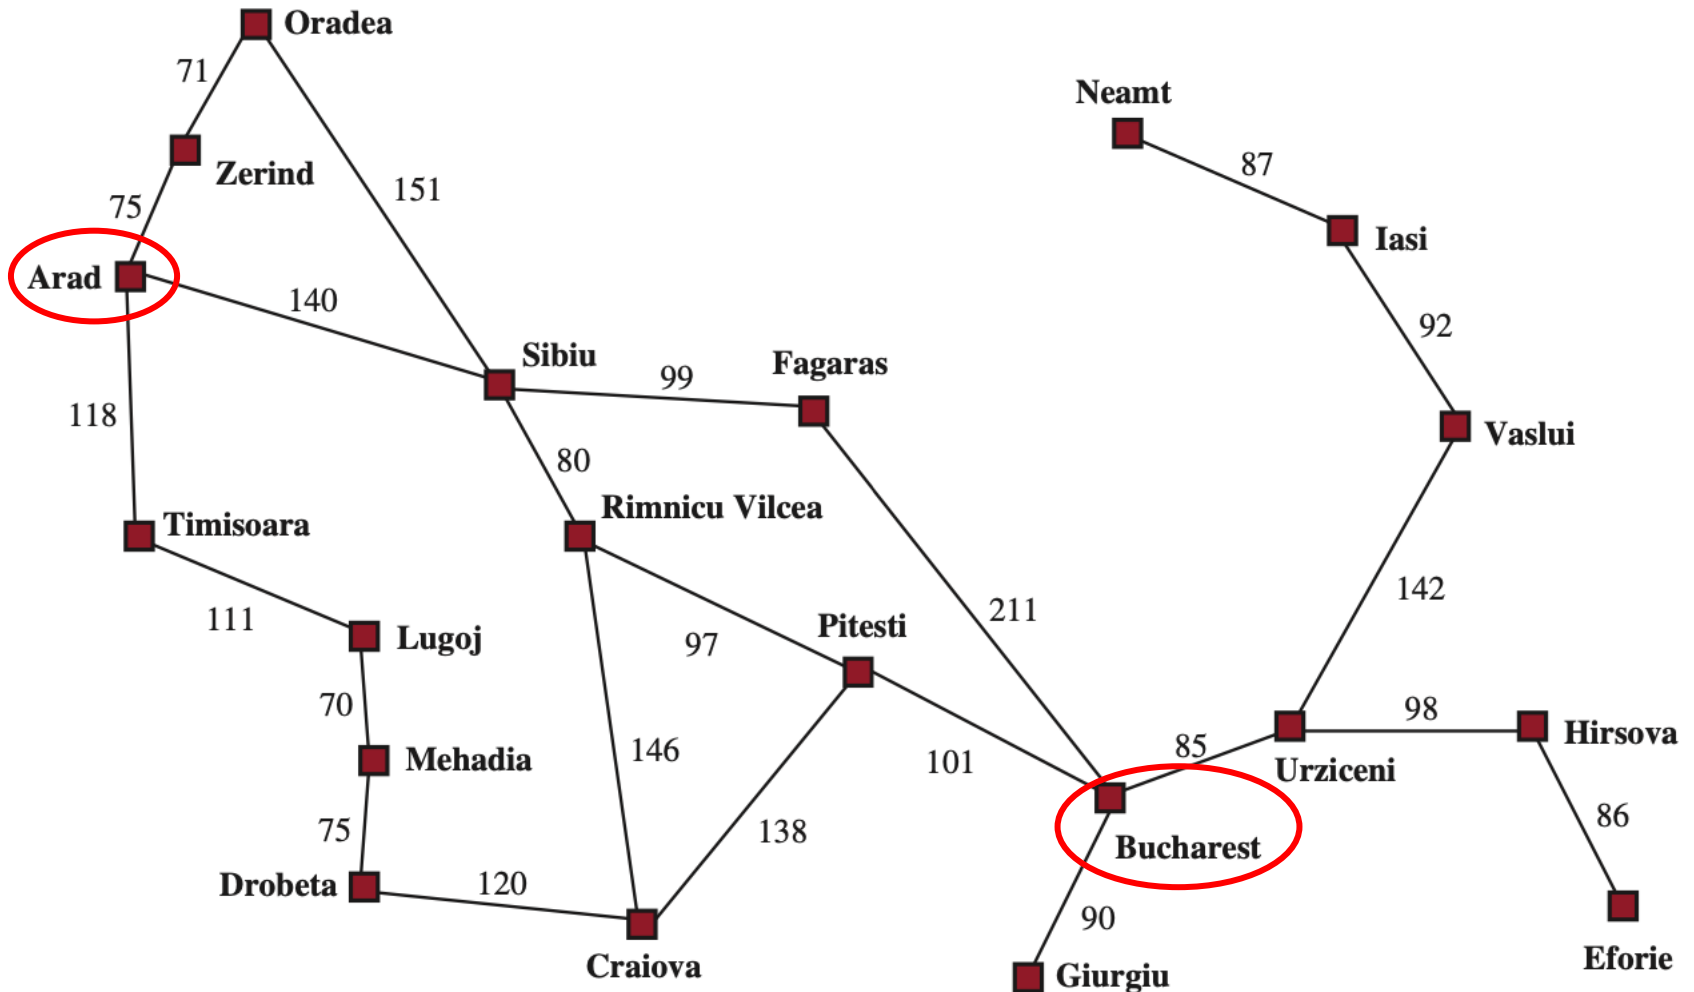
| | | |
|---|---|---|
| 7 | 2 | 4 |
| 5 | | 6 |
| 8 | 3 | 1 |

Start State

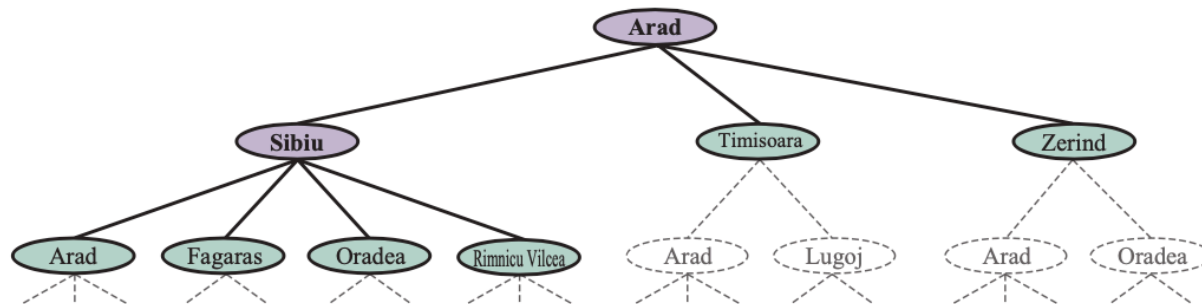
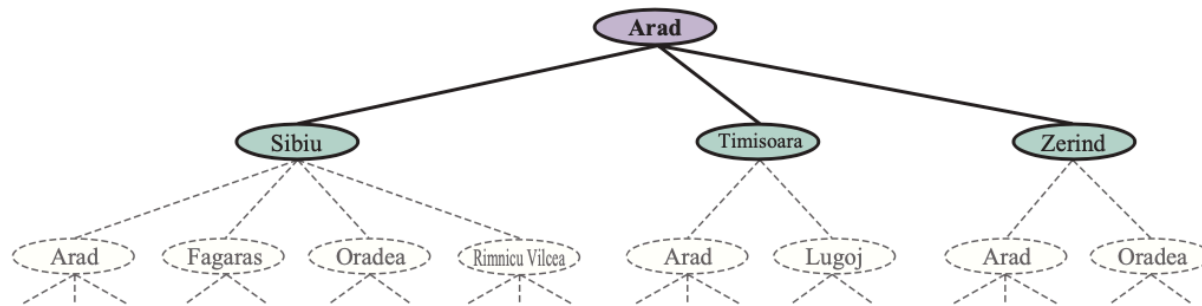
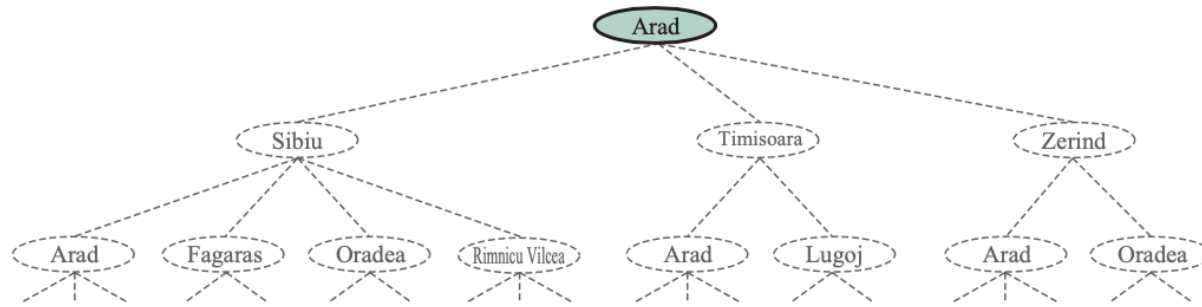
| | | |
|---|---|---|
| | 1 | 2 |
| 3 | 4 | 5 |
| 6 | 7 | 8 |

Goal State

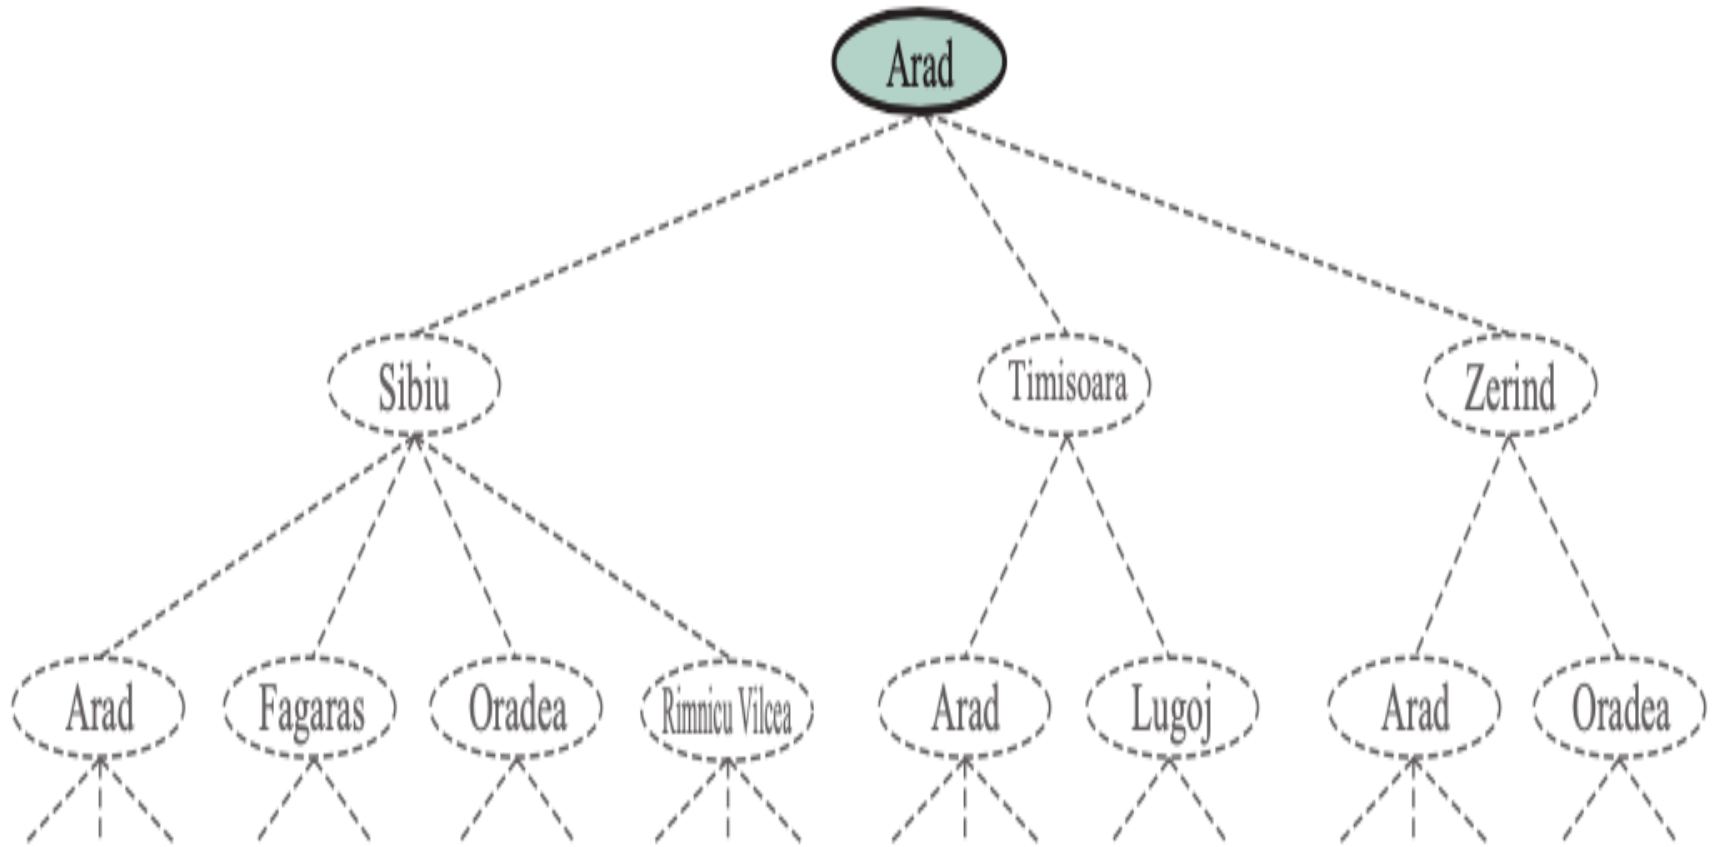
Arad to Bucharest



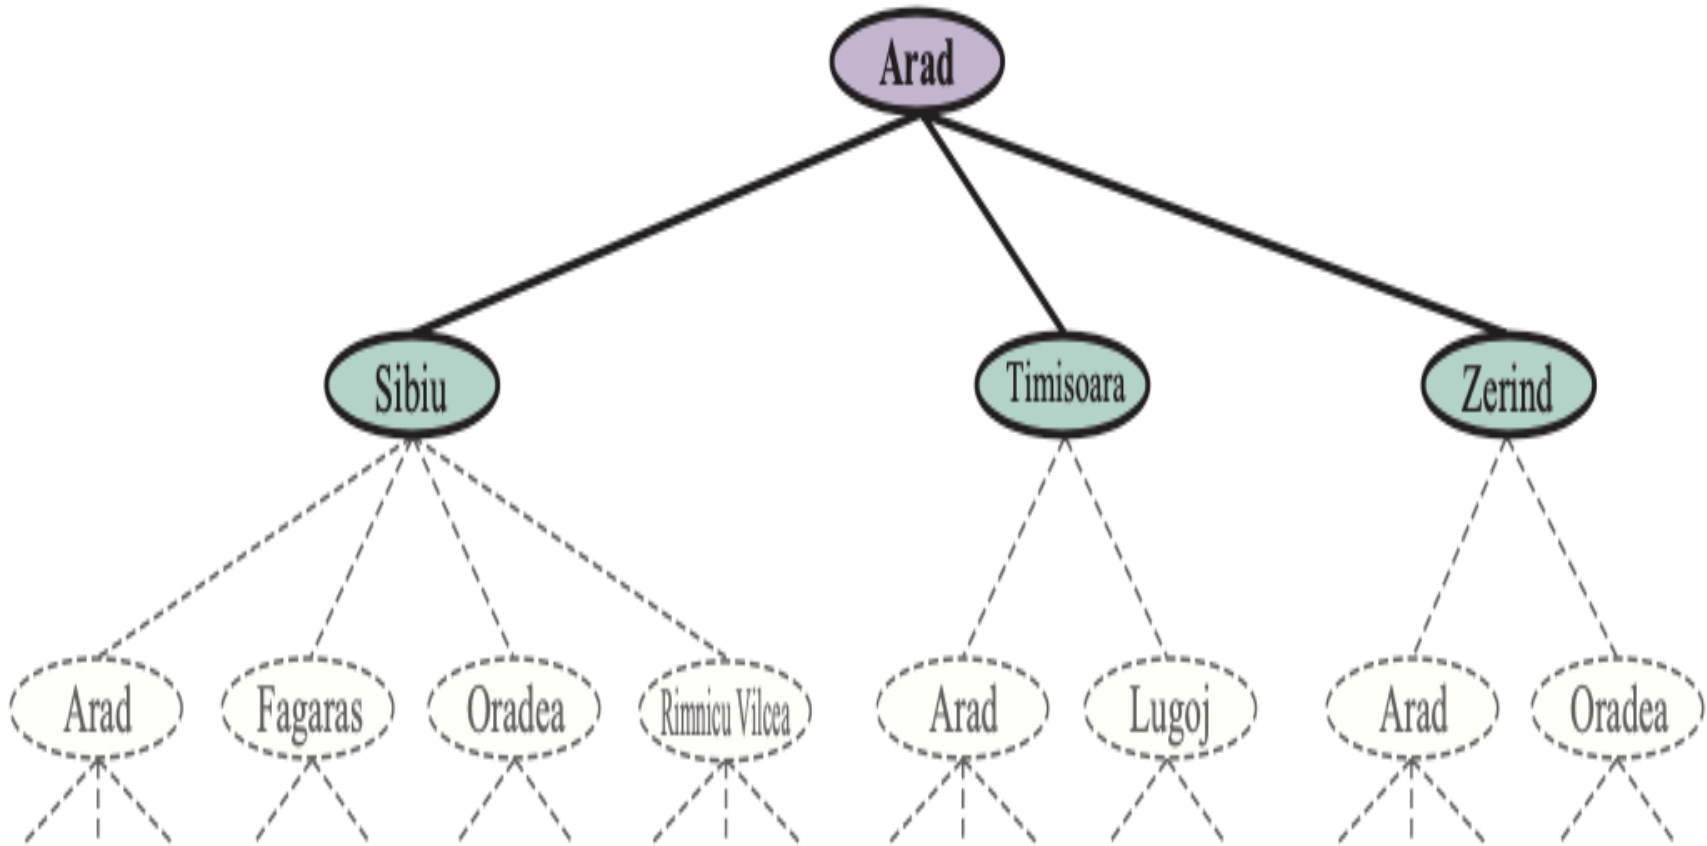
Three partial search trees for finding a route from Arad to Bucharest



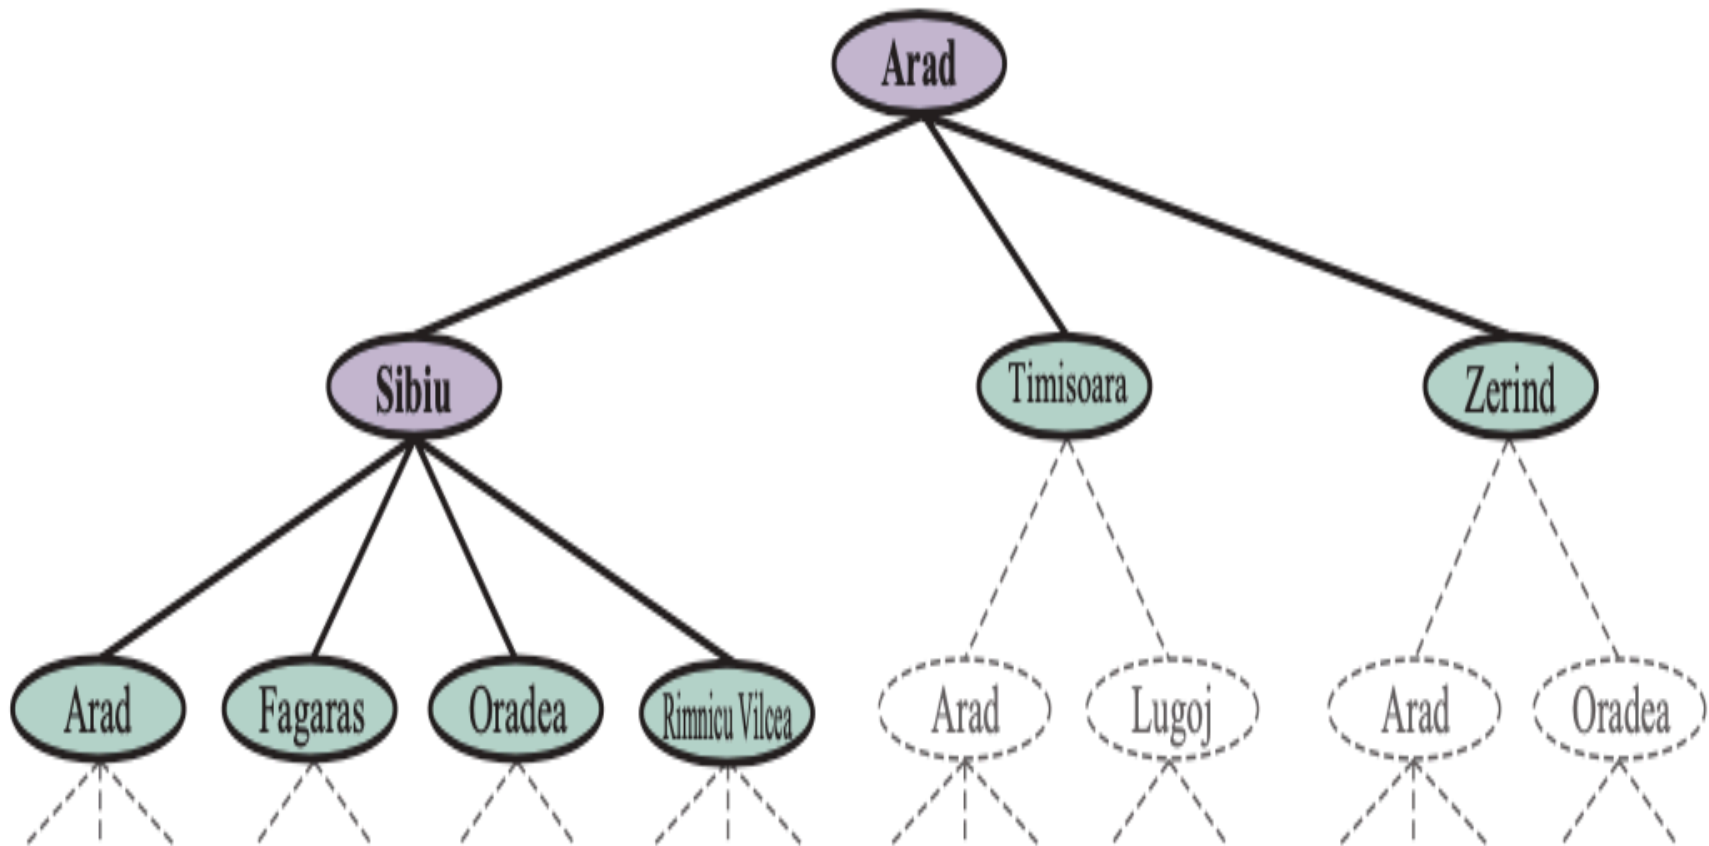
Three partial search trees for finding a route from Arad to Bucharest



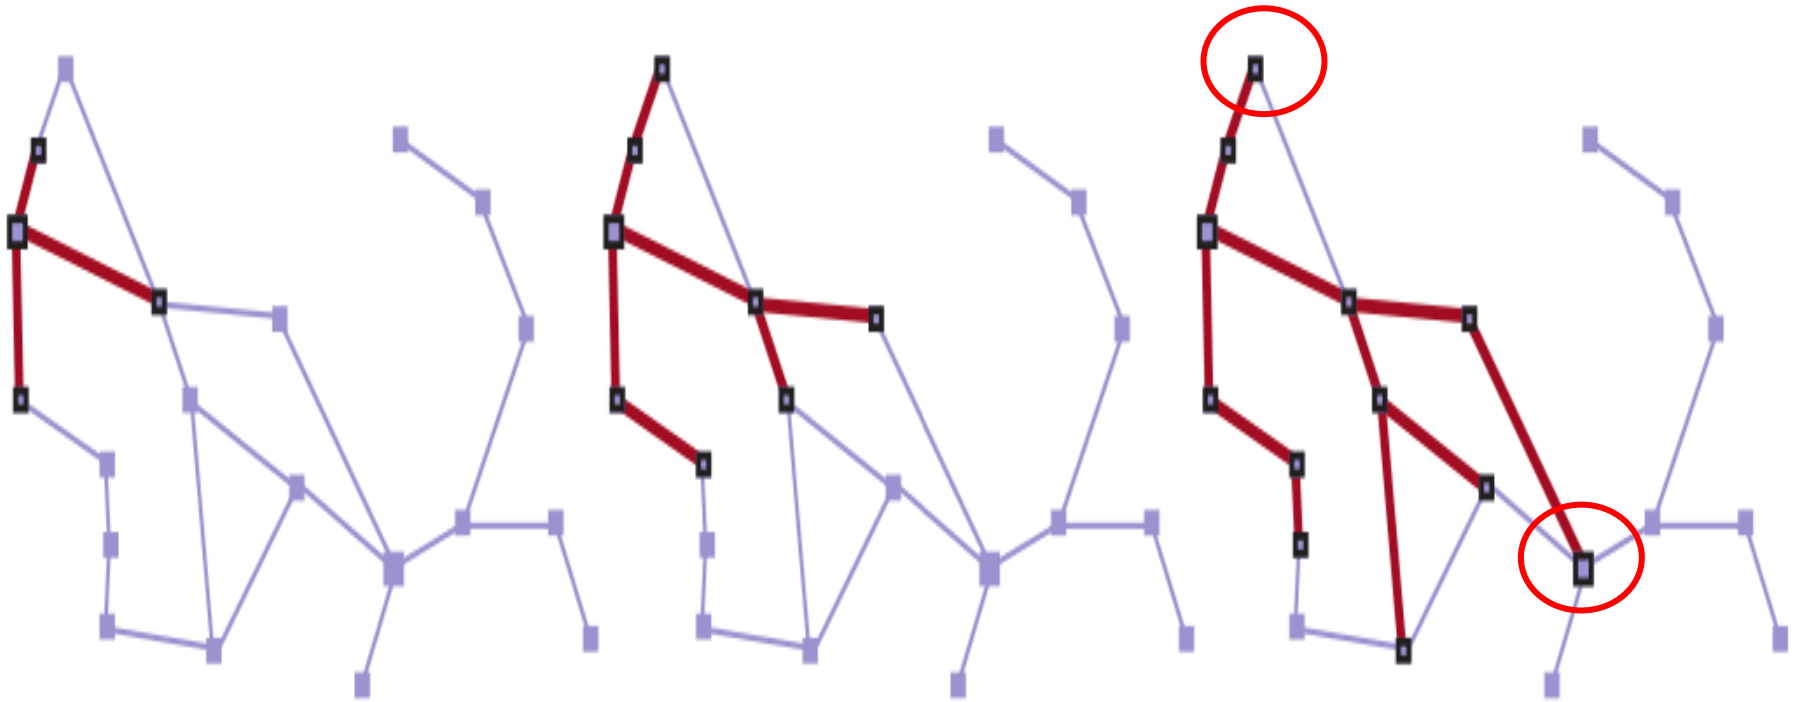
Three partial search trees for finding a route from Arad to Bucharest



Three partial search trees for finding a route from Arad to Bucharest



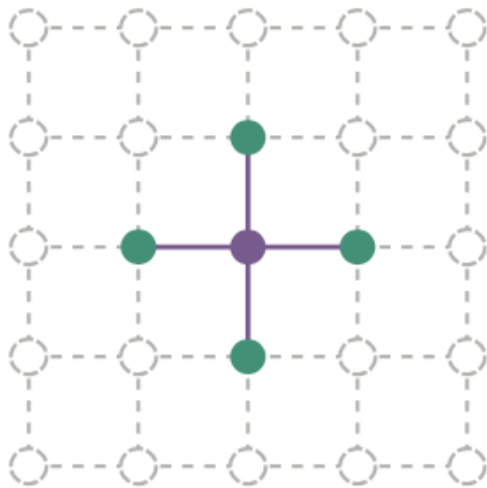
A sequence of search trees generated by a graph search on the Romania problem



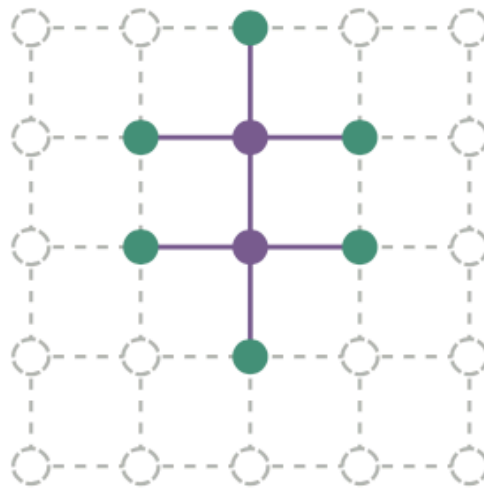
The Separation Property of Graph Search

illustrated on a rectangular-grid problem

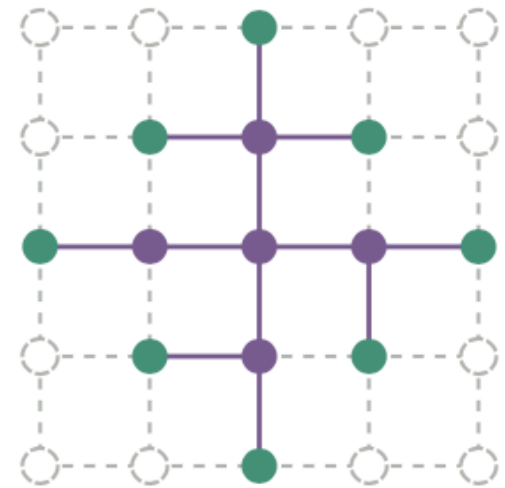
The frontier (green) separates the interior (lavender) from the exterior (faint dashed)



(a)



(b)



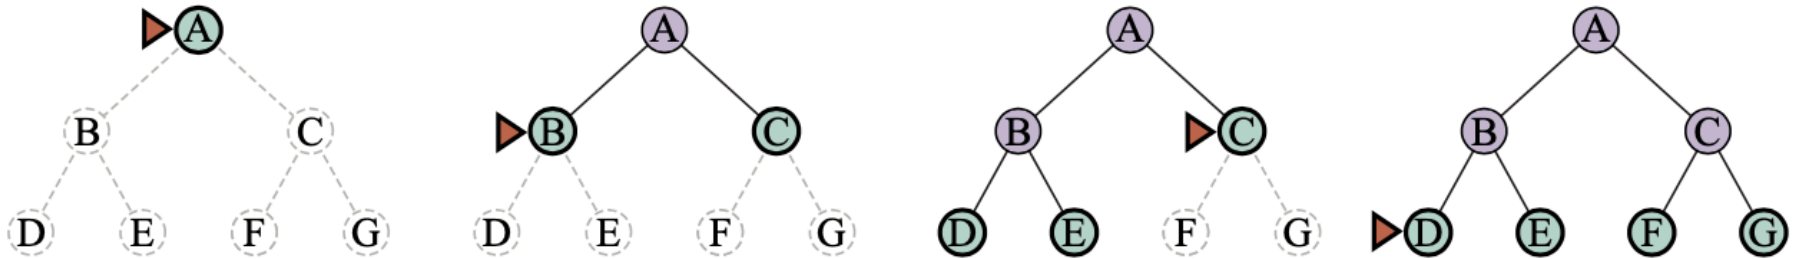
(c)

The Best-First Search (BFS) Algorithm

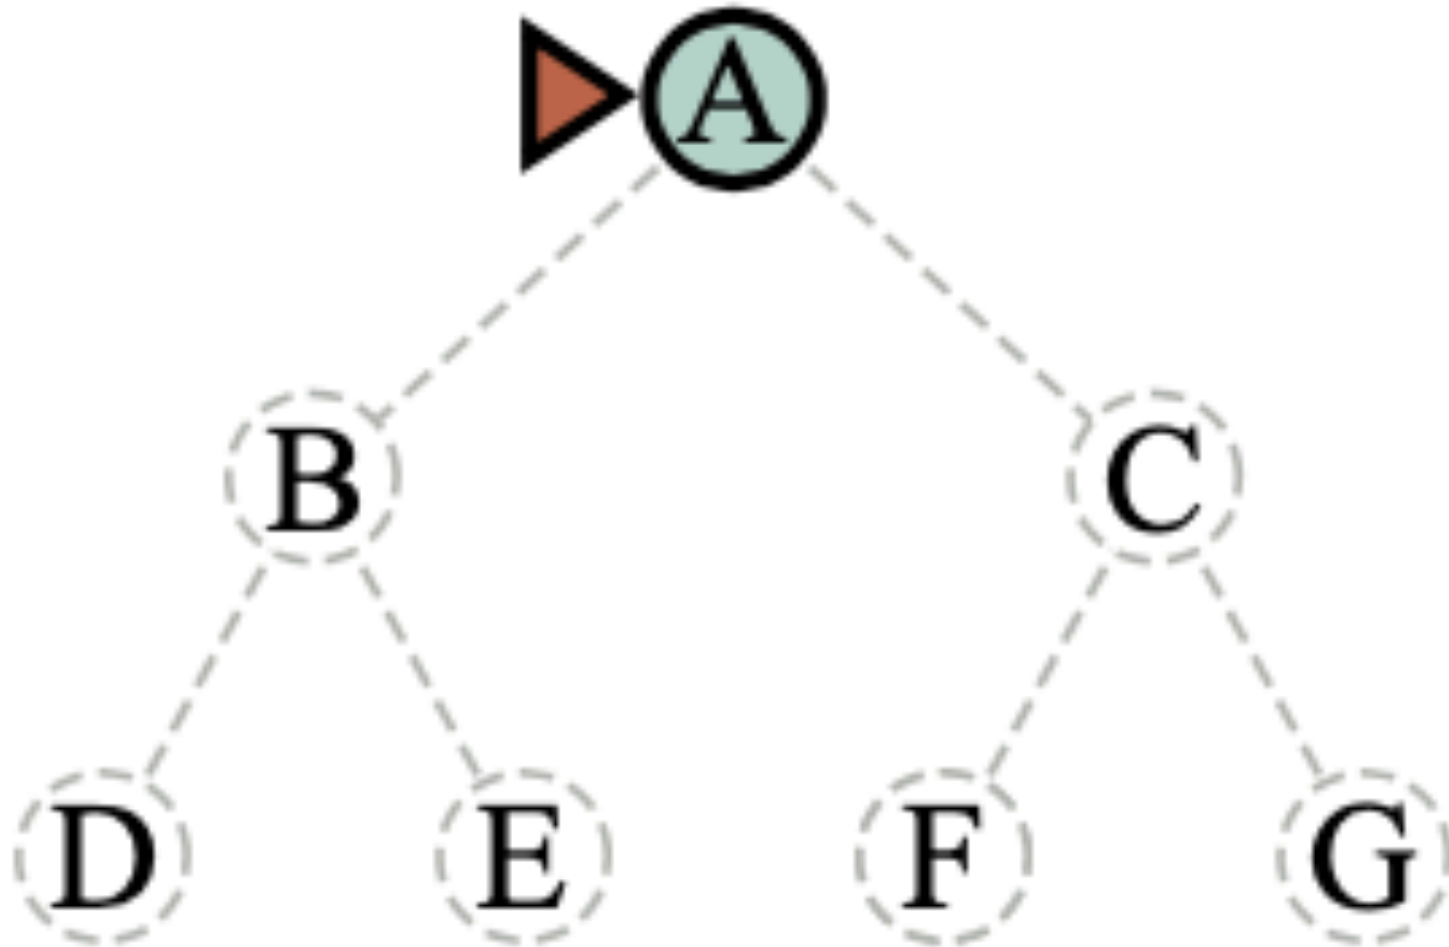
```
function BEST-FIRST-SEARCH(problem, f) returns a solution node or failure  
  node ← NODE(STATE=problem.INITIAL)  
  frontier ← a priority queue ordered by f, with node as an element  
  reached ← a lookup table, with one entry with key problem.INITIAL and value node  
  while not IS-EMPTY(frontier) do  
    node ← POP(frontier)  
    if problem.IS-GOAL(node.STATE) then return node  
    for each child in EXPAND(problem, node) do  
      s ← child.STATE  
      if s is not in reached or child.PATH-COST < reached[s].PATH-COST then  
        reached[s] ← child  
        add child to frontier  
  return failure
```

```
function EXPAND(problem, node) yields nodes  
  s ← node.STATE  
  for each action in problem.ACTIONS(s) do  
    s' ← problem.RESULT(s, action)  
    cost ← node.PATH-COST + problem.ACTION-COST(s, action, s')  
    yield NODE(STATE=s', PARENT=node, ACTION=action, PATH-COST=cost)
```

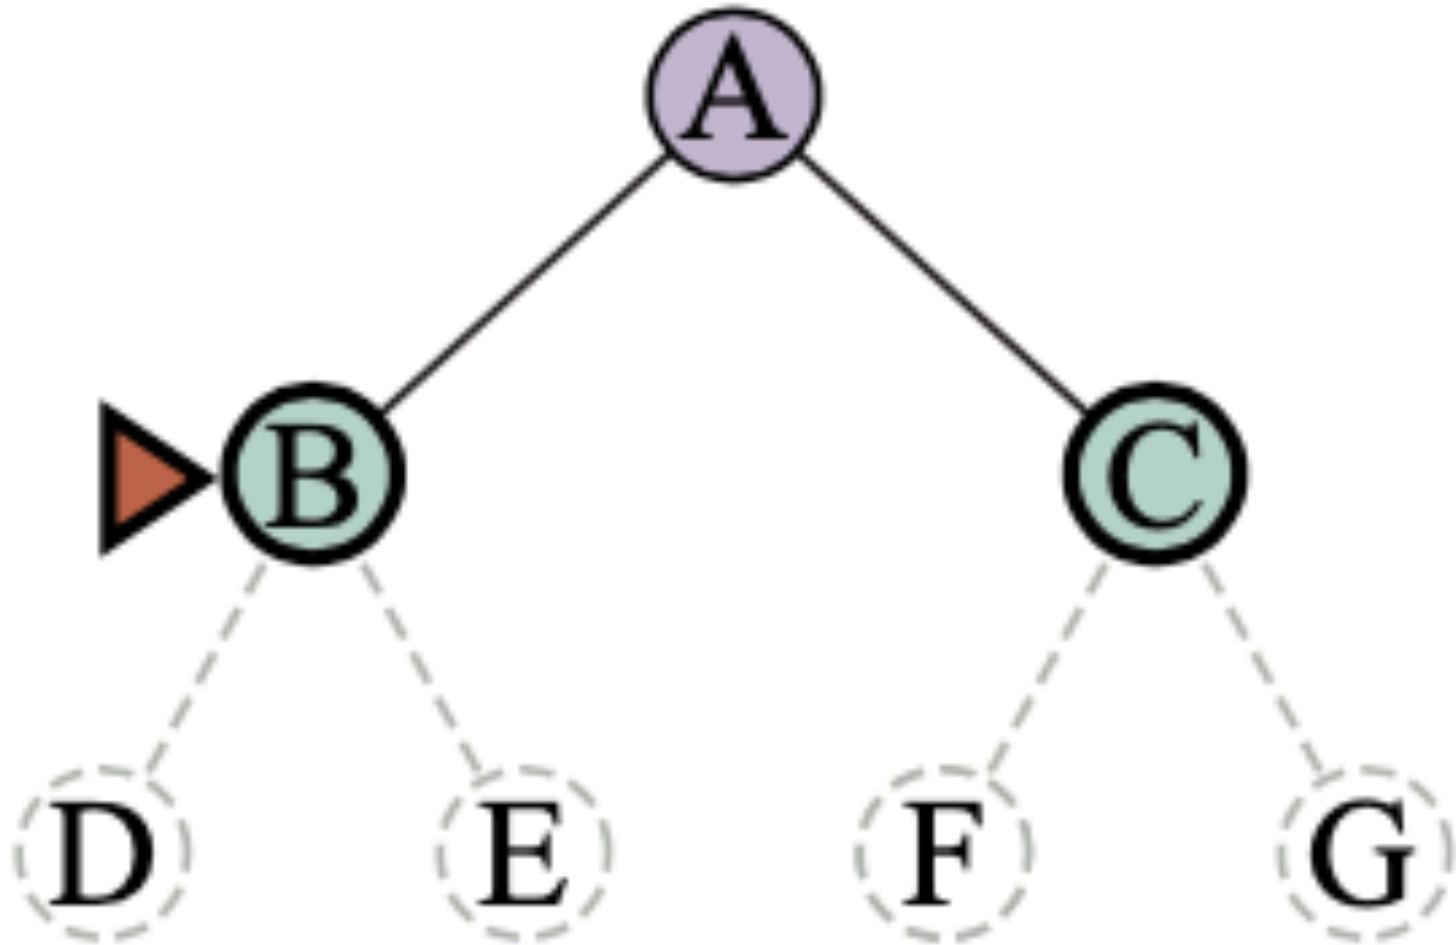
Breadth-First Search on a Simple Binary Tree



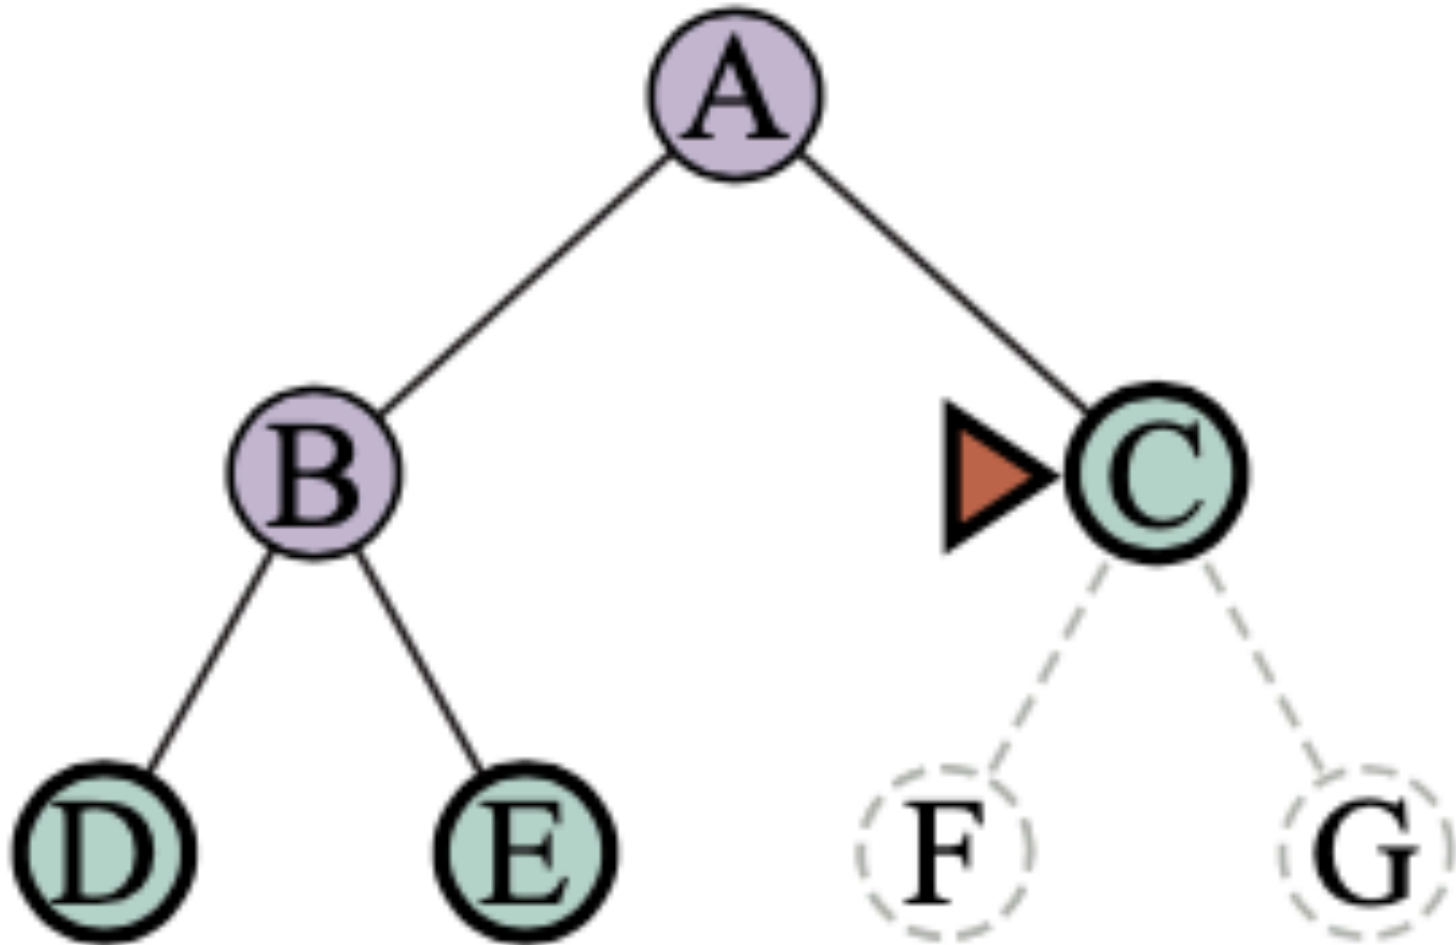
Breadth-First Search on a Simple Binary Tree



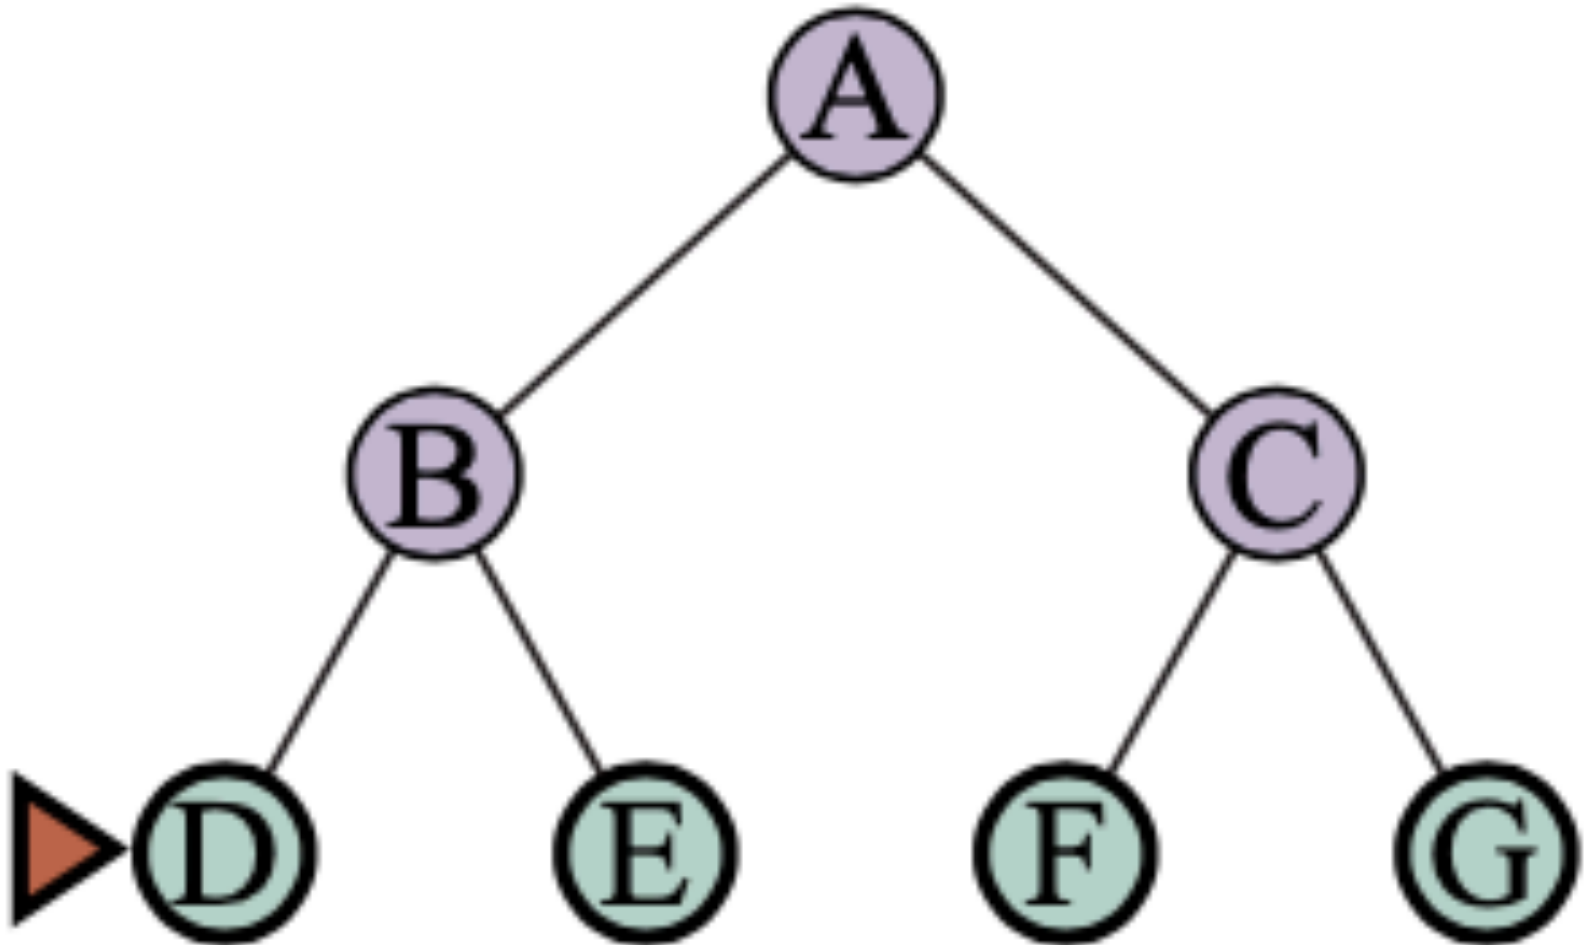
Breadth-First Search on a Simple Binary Tree



Breadth-First Search on a Simple Binary Tree



Breadth-First Search on a Simple Binary Tree



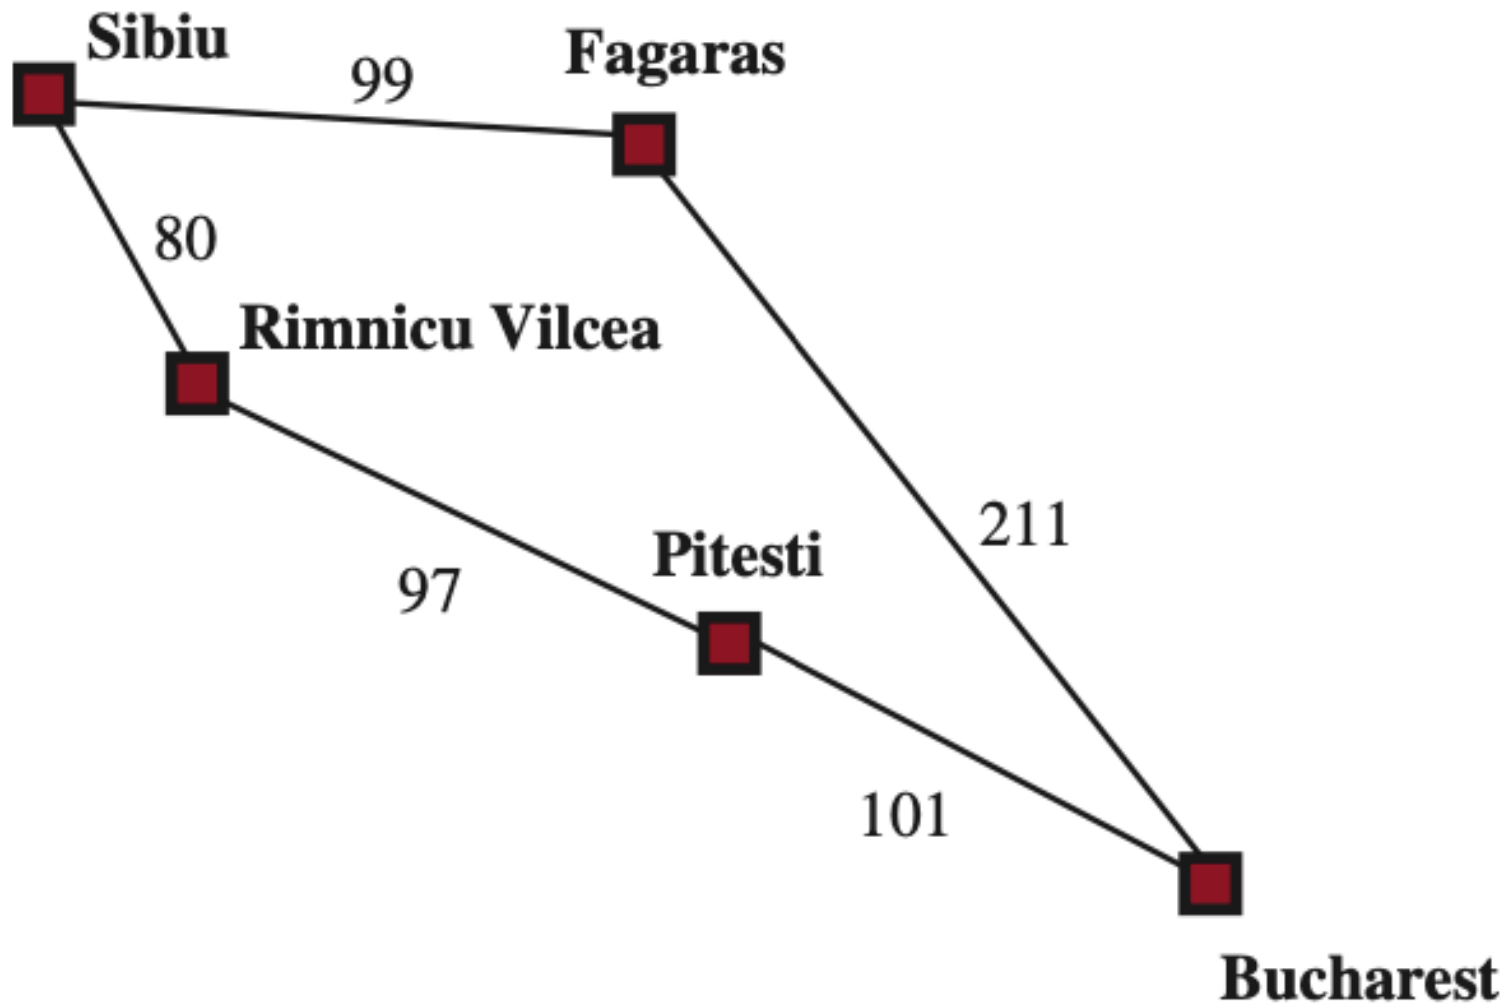
Breadth-First Search and Uniform-Cost Search Algorithms

function BREADTH-FIRST-SEARCH(*problem*) **returns** a solution node or *failure*
 node \leftarrow NODE(*problem*.INITIAL)
 if *problem*.IS-GOAL(*node*.STATE) **then return** *node*
 frontier \leftarrow a FIFO queue, with *node* as an element
 reached \leftarrow {*problem*.INITIAL}
 while not IS-EMPTY(*frontier*) **do**
 node \leftarrow POP(*frontier*)
 for each *child* **in** EXPAND(*problem*, *node*) **do**
 s \leftarrow *child*.STATE
 if *problem*.IS-GOAL(*s*) **then return** *child*
 if *s* is not in *reached* **then**
 add *s* to *reached*
 add *child* to *frontier*
 return *failure*

function UNIFORM-COST-SEARCH(*problem*) **returns** a solution node, or *failure*
 return BEST-FIRST-SEARCH(*problem*, PATH-COST)

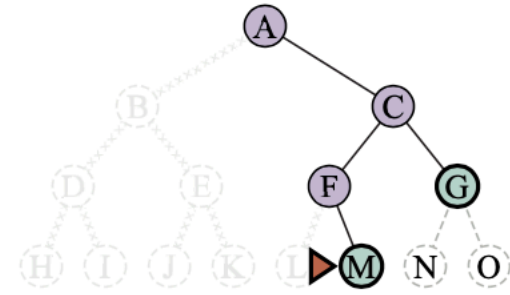
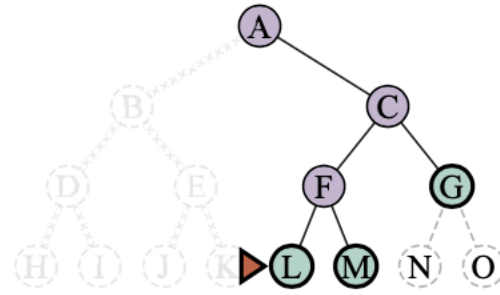
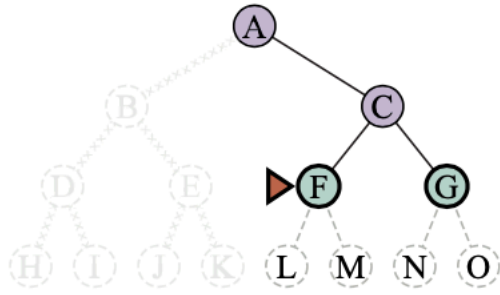
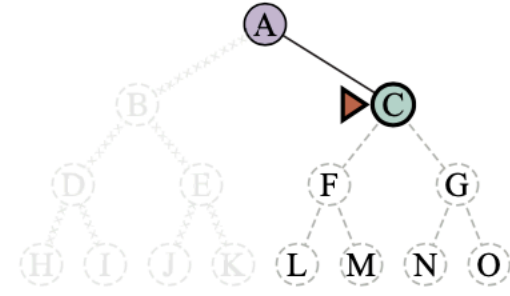
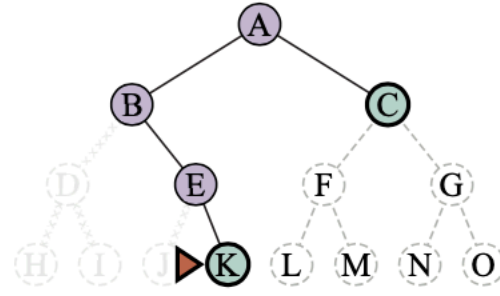
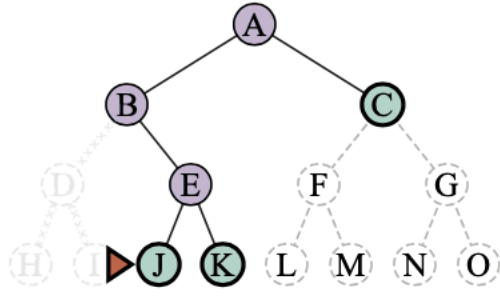
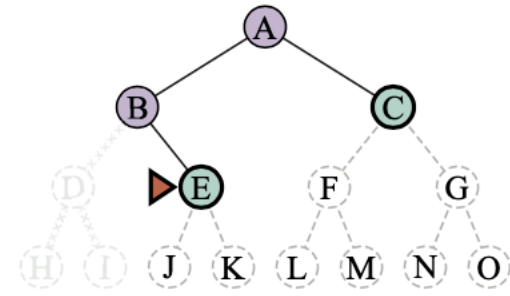
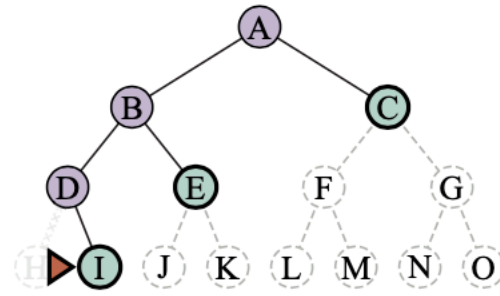
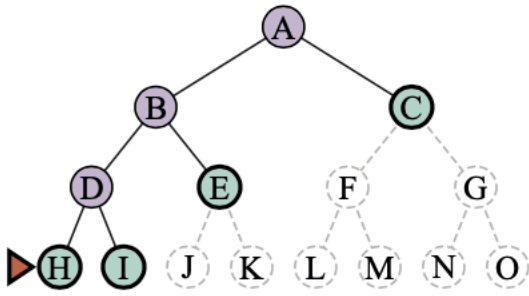
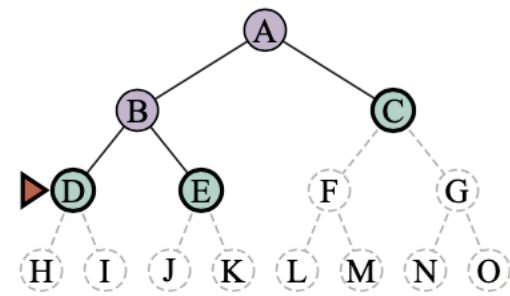
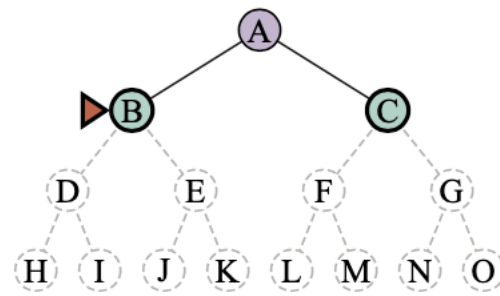
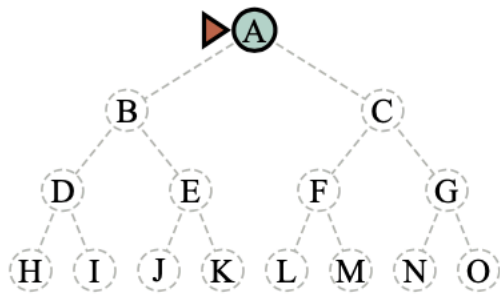
Part of the Romania State Space

Uniform-Cost Search



Depth-First Search (DFS)

Depth-First Search (DFS)

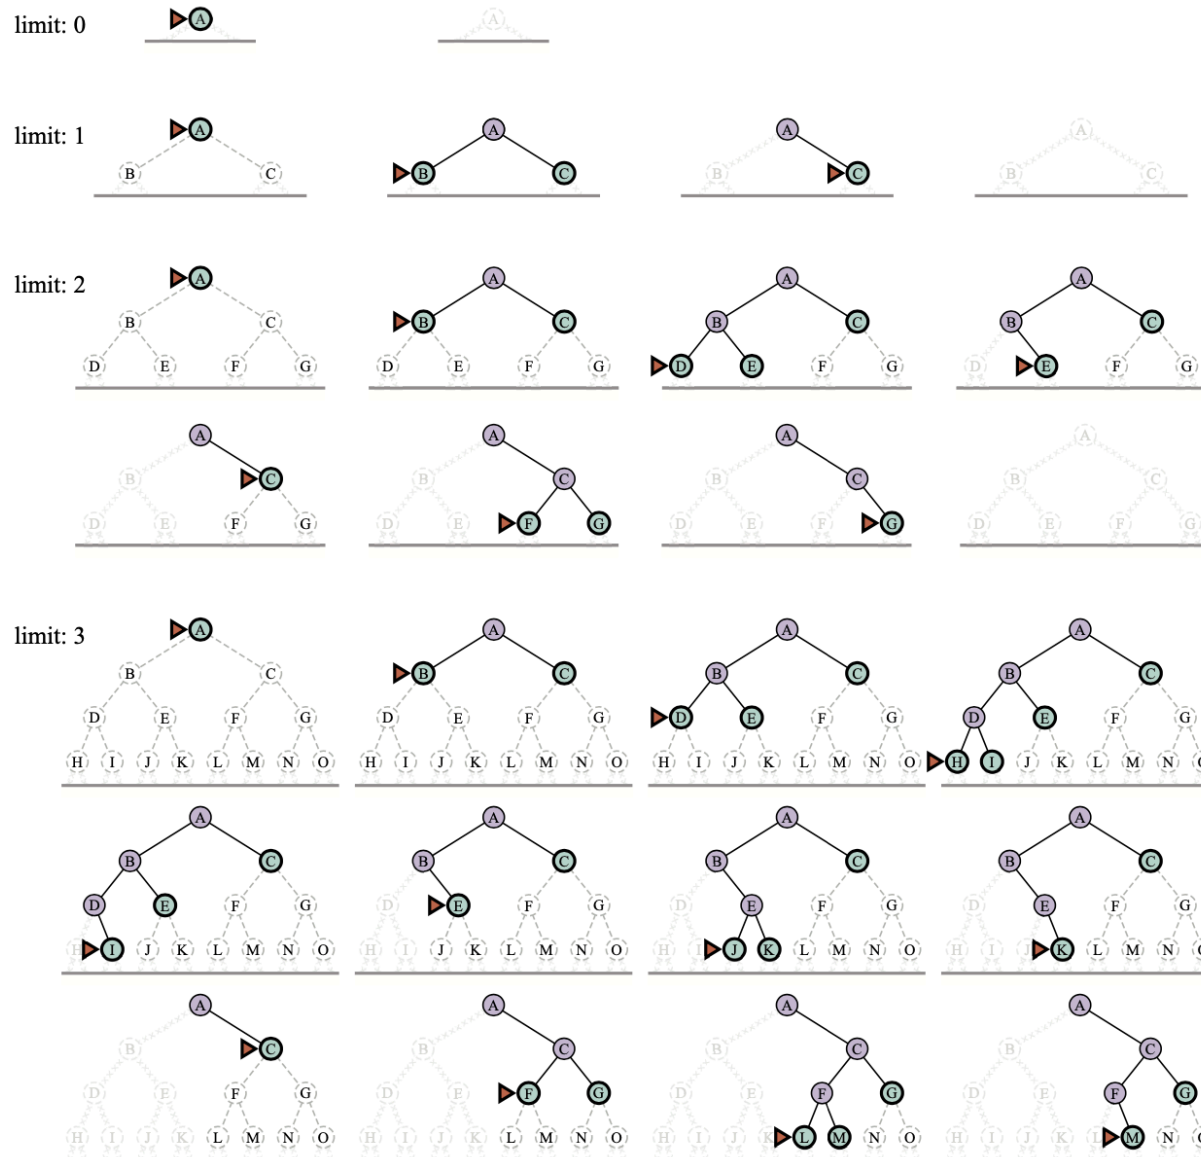


Iterative deepening and depth-limited tree-like search

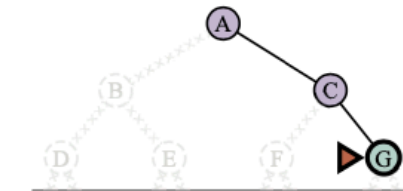
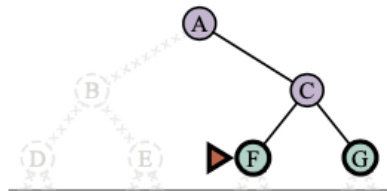
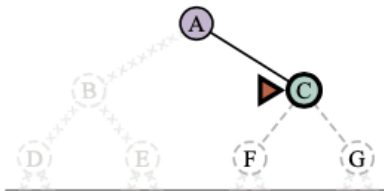
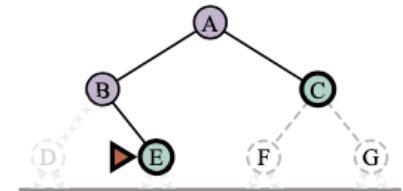
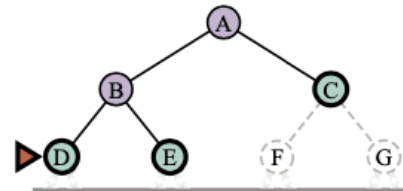
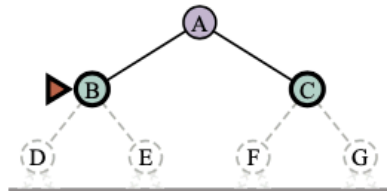
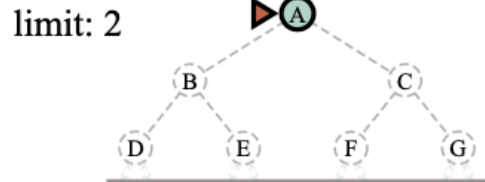
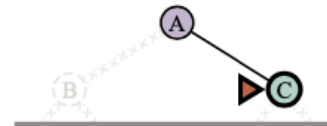
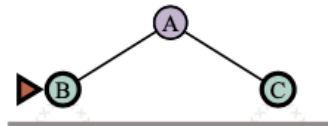
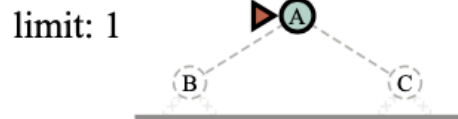
function ITERATIVE-DEEPENING-SEARCH(*problem*) **returns** a solution node or *failure*
for *depth* = 0 **to** ∞ **do**
 result \leftarrow DEPTH-LIMITED-SEARCH(*problem*, *depth*)
 if *result* \neq *cutoff* **then return** *result*

function DEPTH-LIMITED-SEARCH(*problem*, ℓ) **returns** a node or *failure* or *cutoff*
frontier \leftarrow a LIFO queue (stack) with NODE(*problem*.INITIAL) as an element
result \leftarrow *failure*
while not IS-EMPTY(*frontier*) **do**
 node \leftarrow POP(*frontier*)
 if *problem*.IS-GOAL(*node*.STATE) **then return** *node*
 if DEPTH(*node*) > ℓ **then**
 result \leftarrow *cutoff*
 else if not IS-CYCLE(*node*) **do**
 for each *child* **in** EXPAND(*problem*, *node*) **do**
 add *child* to *frontier*
return *result*

Four iterations of iterative deepening search

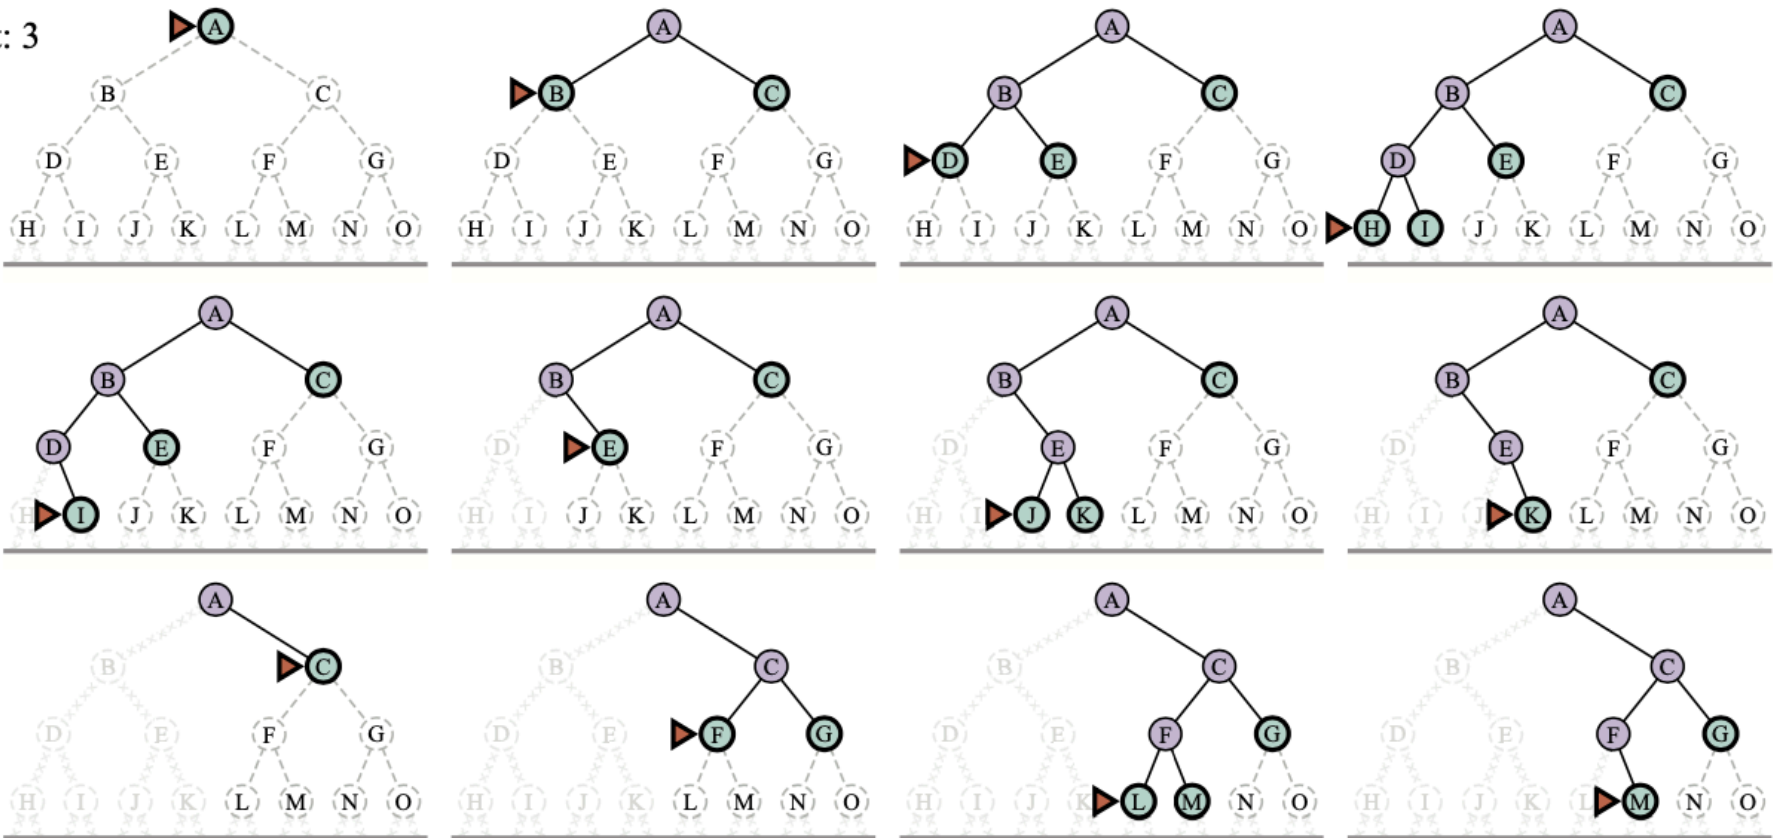


Four iterations of iterative deepening search



Four iterations of iterative deepening search

limit: 3



Bidirectional Best-First Search

keeps two frontiers and two tables of reached states

```
function BIBF-SEARCH(problemF, fF, problemB, fB) returns a solution node, or failure  
  nodeF ← NODE(problemF.INITIAL) // Node for a start state  
  nodeB ← NODE(problemB.INITIAL) // Node for a goal state  
  frontierF ← a priority queue ordered by fF, with nodeF as an element  
  frontierB ← a priority queue ordered by fB, with nodeB as an element  
  reachedF ← a lookup table, with one key nodeF.STATE and value nodeF  
  reachedB ← a lookup table, with one key nodeB.STATE and value nodeB  
  solution ← failure  
  while not TERMINATED(solution, frontierF, frontierB) do  
    if fF(TOP(frontierF)) < fB(TOP(frontierB)) then  
      solution ← PROCEED(F, problemF, frontierF, reachedF, reachedB, solution)  
    else solution ← PROCEED(B, problemB, frontierB, reachedB, reachedF, solution)  
  return solution
```

Bidirectional Best-First Search

keeps two frontiers and two tables of reached states

```
function PROCEED(dir, problem, frontier, reached, reached2, solution) returns a solution
    // Expand node on frontier; check against the other frontier in reached2.
    // The variable “dir” is the direction: either F for forward or B for backward.
    node ← POP(frontier)
    for each child in EXPAND(problem, node) do
        s ← child.STATE
        if s not in reached or PATH-COST(child) < PATH-COST(reached[s]) then
            reached[s] ← child
            add child to frontier
            if s is in reached2 then
                solution2 ← JOIN-NODES(dir, child, reached2[s])
                if PATH-COST(solution2) < PATH-COST(solution) then
                    solution ← solution2
    return solution
```

Evaluation of search algorithms

| Criterion | Breadth-First | Uniform-Cost | Depth-First | Depth-Limited | Iterative Deepening | Bidirectional (if applicable) |
|---------------|------------------|---------------------------------------|-------------|---------------|---------------------|-------------------------------|
| Complete? | Yes ¹ | Yes ^{1,2} | No | No | Yes ¹ | Yes ^{1,4} |
| Optimal cost? | Yes ³ | Yes | No | No | Yes ³ | Yes ^{3,4} |
| Time | $O(b^d)$ | $O(b^{1+\lceil C^*/\epsilon \rceil})$ | $O(b^m)$ | $O(b^\ell)$ | $O(b^d)$ | $O(b^{d/2})$ |
| Space | $O(b^d)$ | $O(b^{1+\lceil C^*/\epsilon \rceil})$ | $O(bm)$ | $O(b\ell)$ | $O(bd)$ | $O(b^{d/2})$ |

b is the branching factor; m is the maximum depth of the search tree; d is the depth of the shallowest solution, or is m when there is no solution; ℓ is the depth limit

Values of $hSLD$

—straight-line distances to Bucharest.

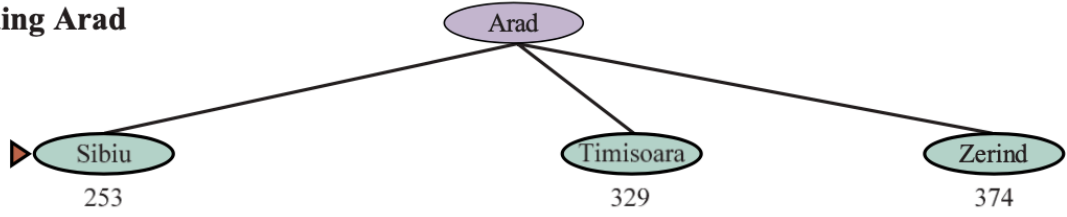
| | | | |
|------------------|-----|-----------------------|-----|
| Arad | 366 | Mehadia | 241 |
| Bucharest | 0 | Neamt | 234 |
| Craiova | 160 | Oradea | 380 |
| Drobeta | 242 | Pitesti | 100 |
| Eforie | 161 | Rimnicu Vilcea | 193 |
| Fagaras | 176 | Sibiu | 253 |
| Giurgiu | 77 | Timisoara | 329 |
| Hirsova | 151 | Urziceni | 80 |
| Iasi | 226 | Vaslui | 199 |
| Lugoj | 244 | Zerind | 374 |

A* search

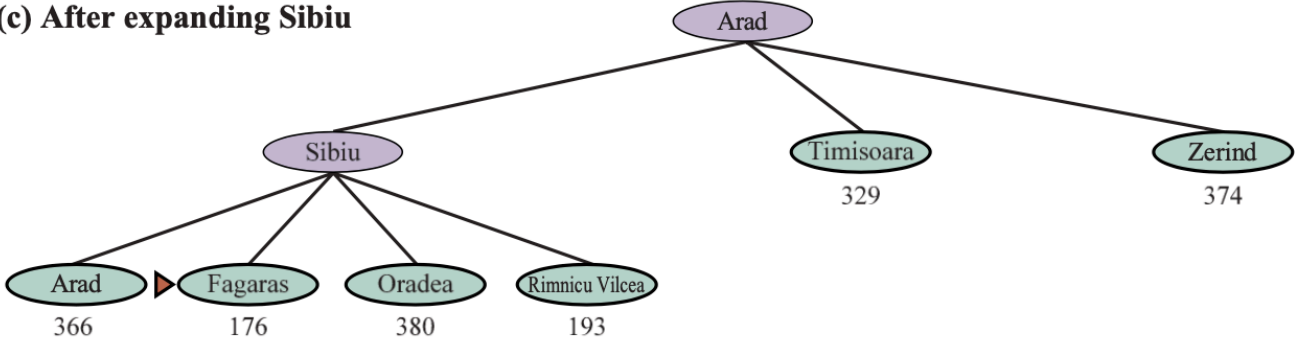
(a) The initial state



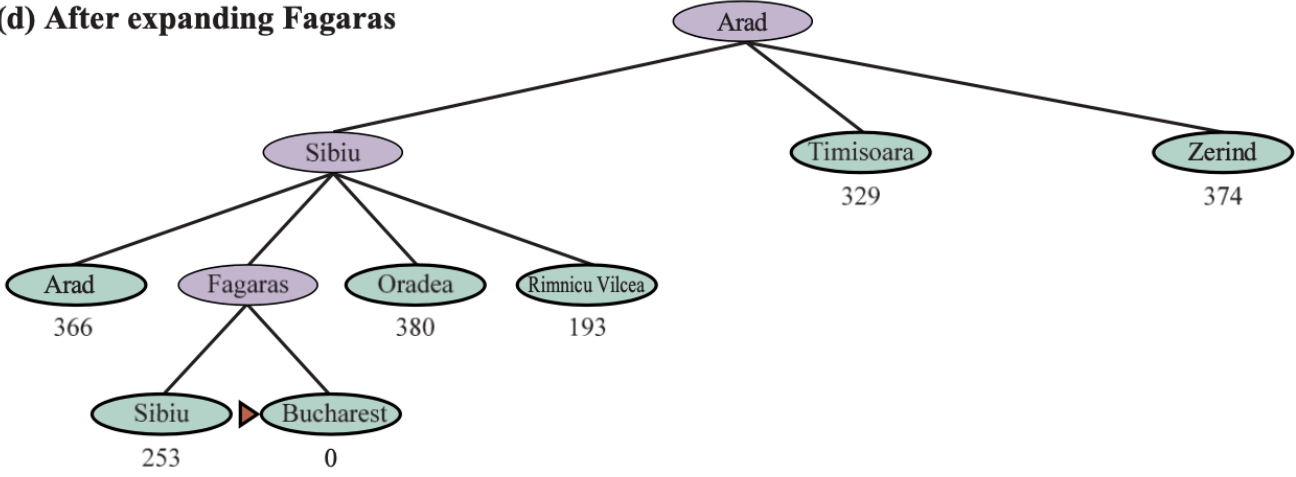
(b) After expanding Arad



(c) After expanding Sibiu



(d) After expanding Fagaras



A* search

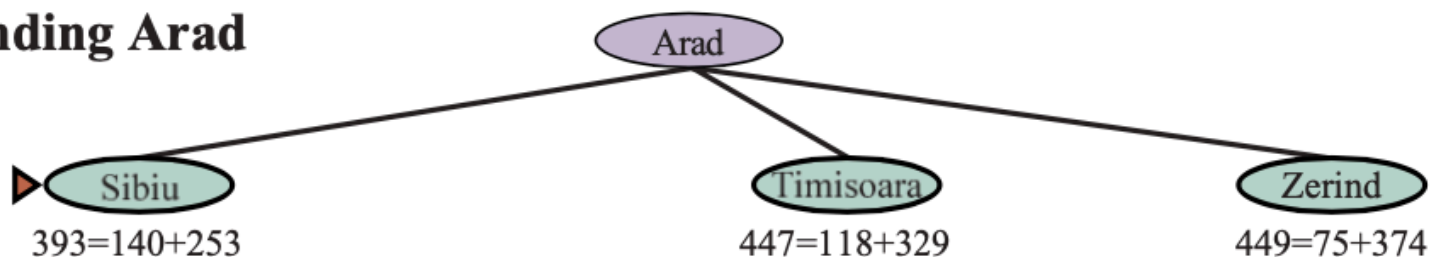
Nodes are labeled with $f = g + h$.

The h values are the Straight-Line Distances heuristic h_{SLD}

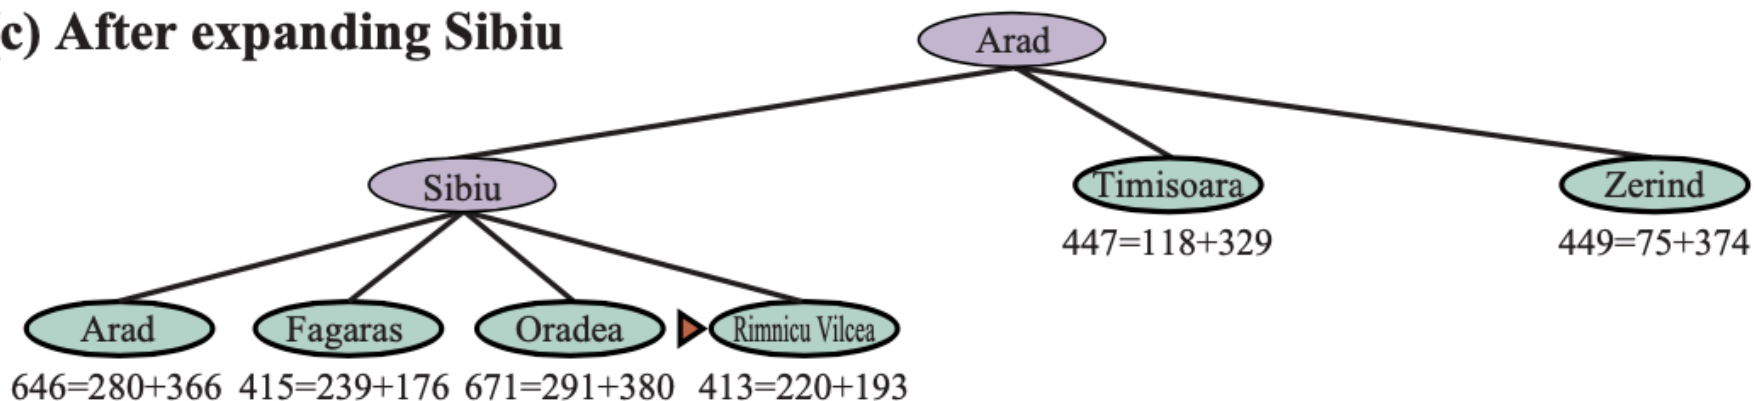
(a) The initial state



(b) After expanding Arad



(c) After expanding Sibiu

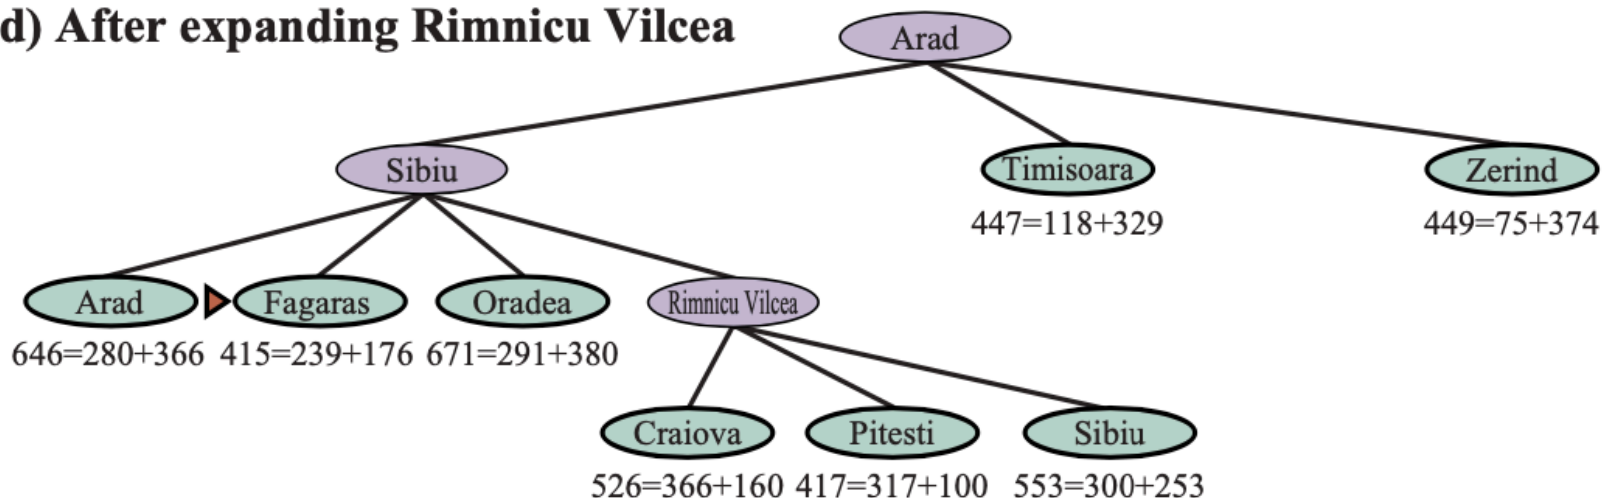


A* search

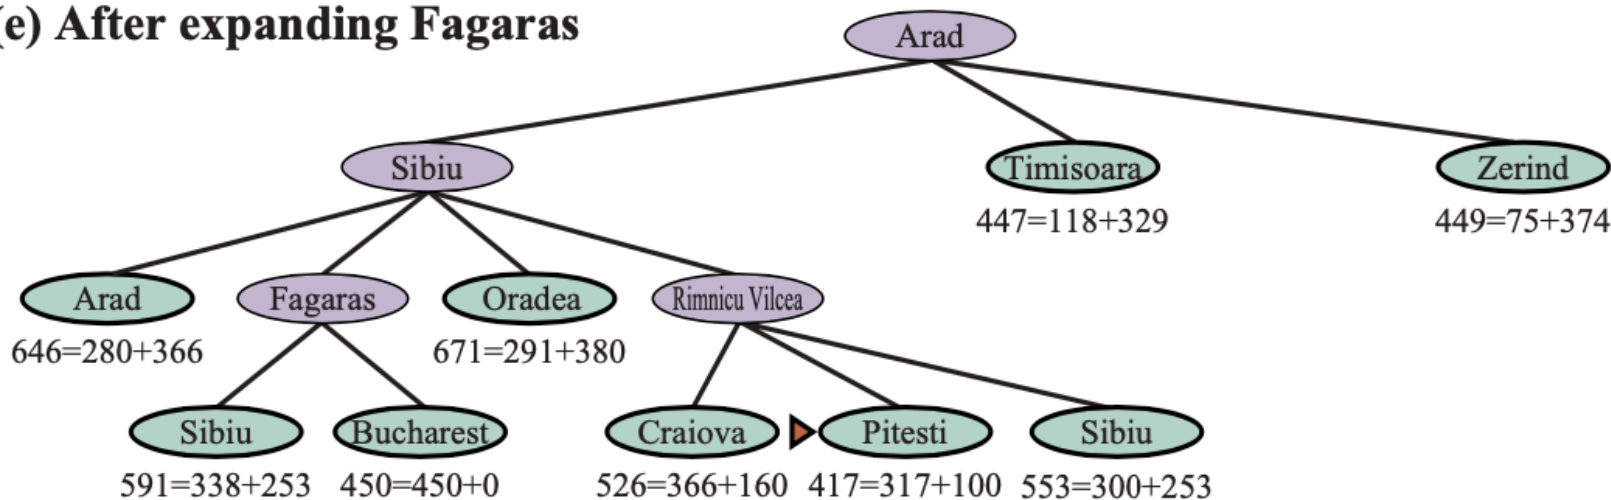
Nodes are labeled with $f = g + h$.

The h values are the Straight-Line Distances heuristic h_{SLD}

(d) After expanding Rimnicu Vilcea



(e) After expanding Fagaras

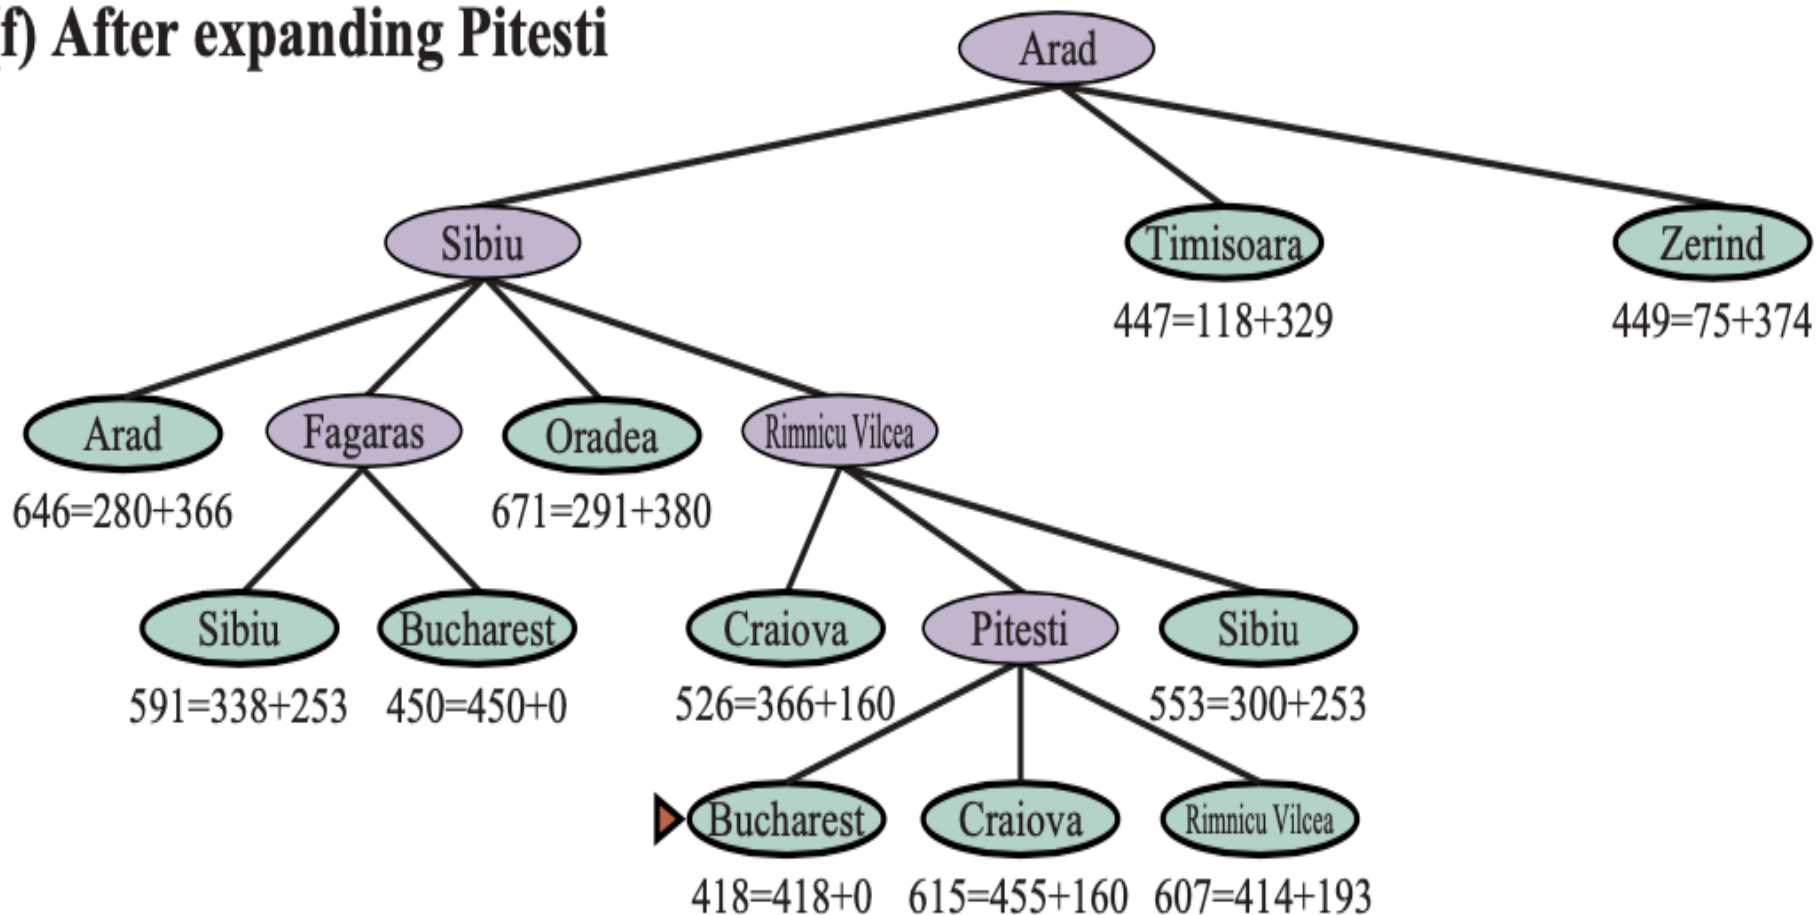


A* search

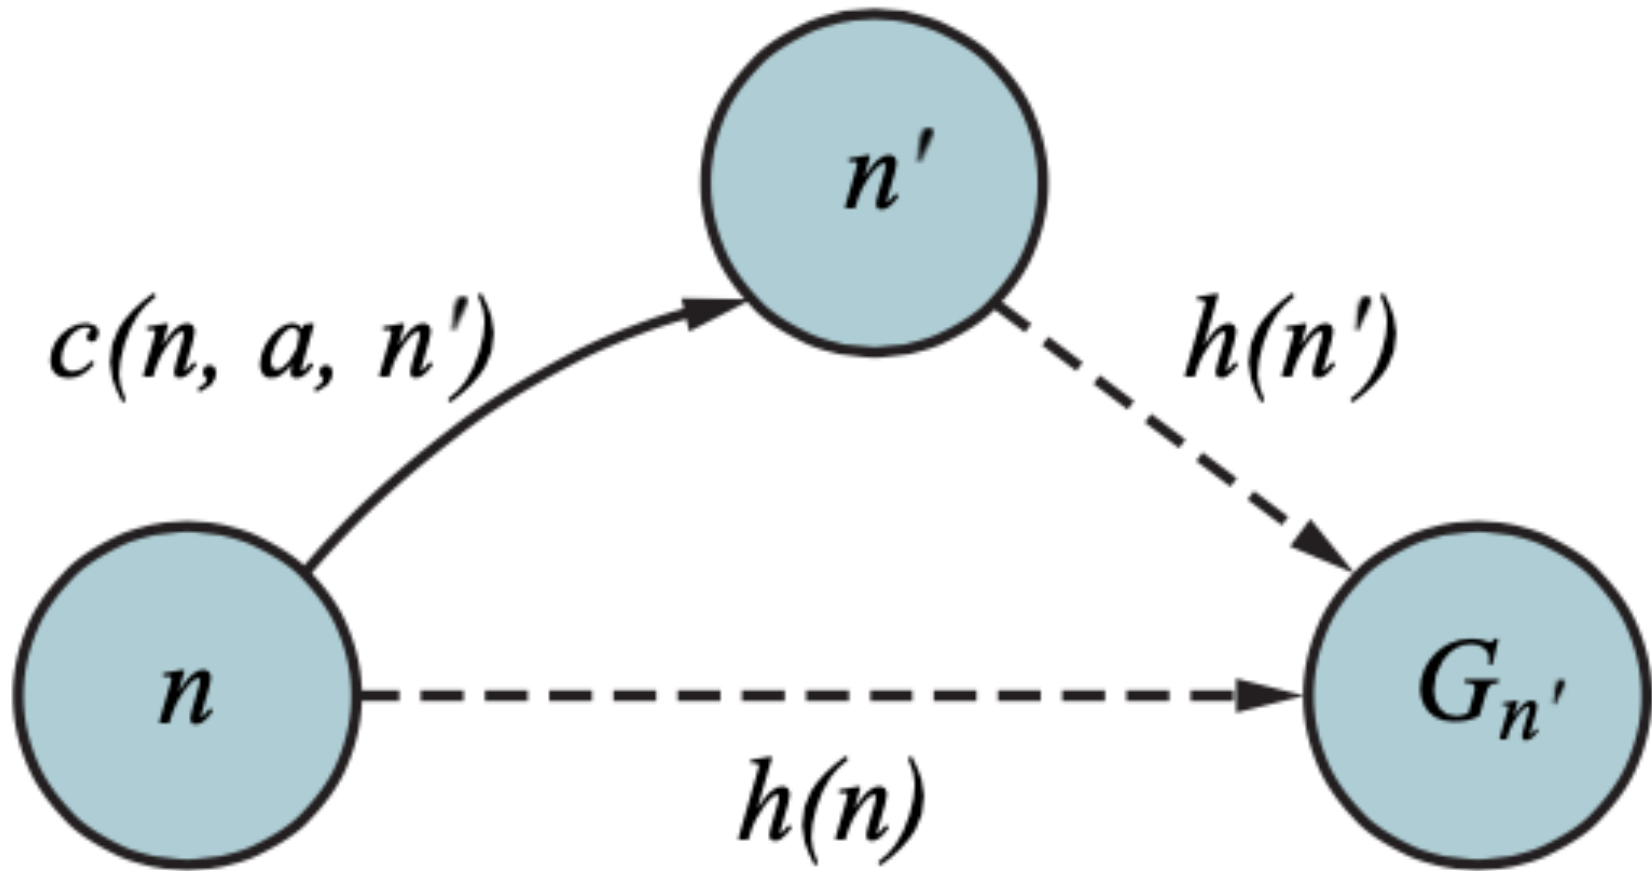
Nodes are labeled with $f = g + h$.

The h values are the Straight-Line Distances heuristic h_{SLD}

(f) After expanding Pitesti

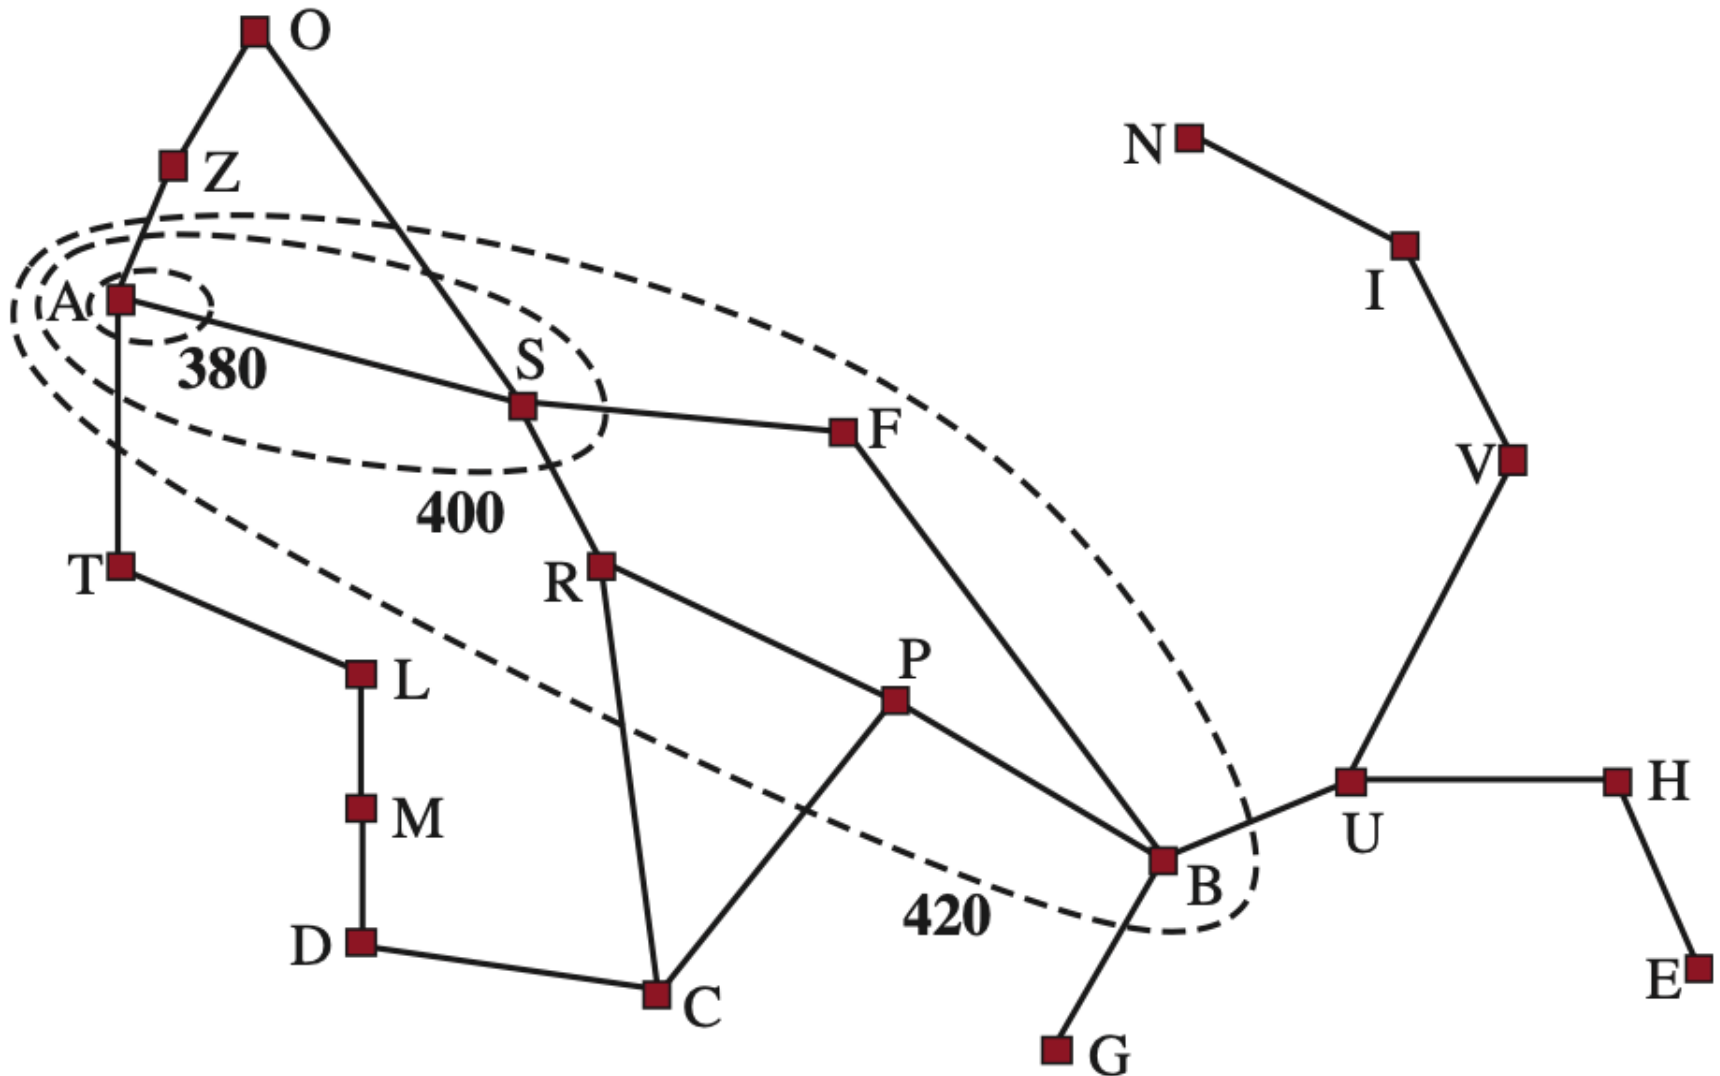


Triangle Inequality



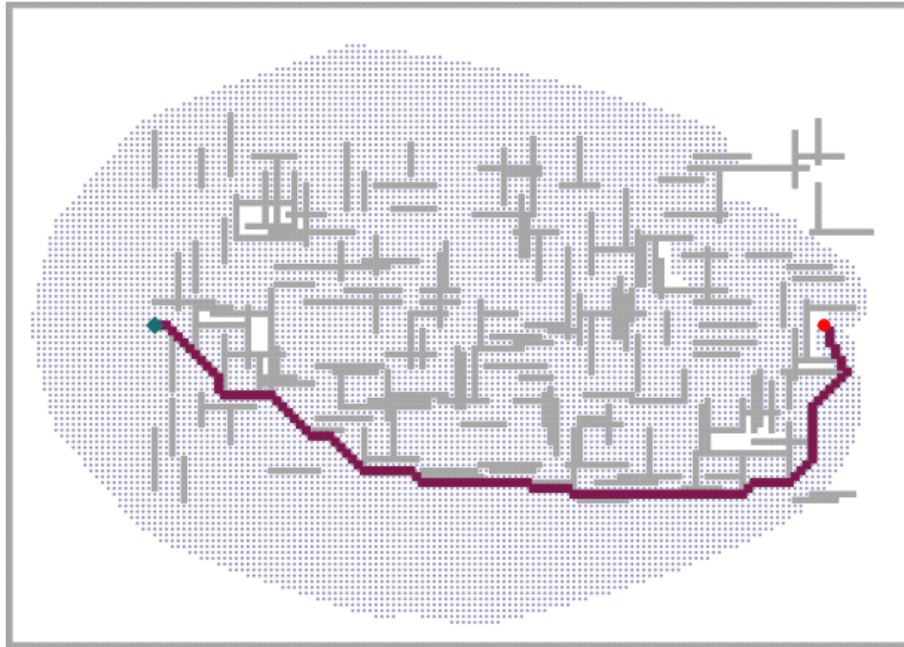
If the heuristic h is consistent, then the single number $h(n)$ will be less than the sum of the cost $c(n, a, a')$ of the action from n to n' plus the heuristic estimate $h(n')$.

Map of Romania showing contours at $f = 380$, $f = 400$, and $f = 420$, with Arad as the start state

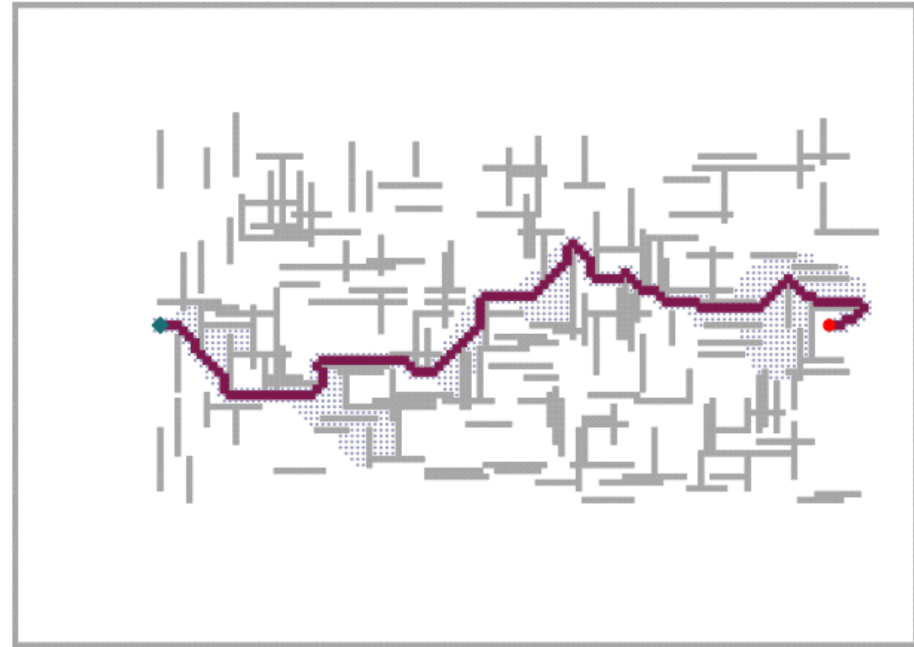


(a) A* Search

(b) Weighted A* Search



(a)



(b)

The gray bars are obstacles, the purple line is the path from the green start to red goal, and the small dots are states that were reached by each search.

On this particular problem, weighted A* explores 7 times fewer states and finds a path that is 5% more costly.

Recursive Best-First Search (RBFS)

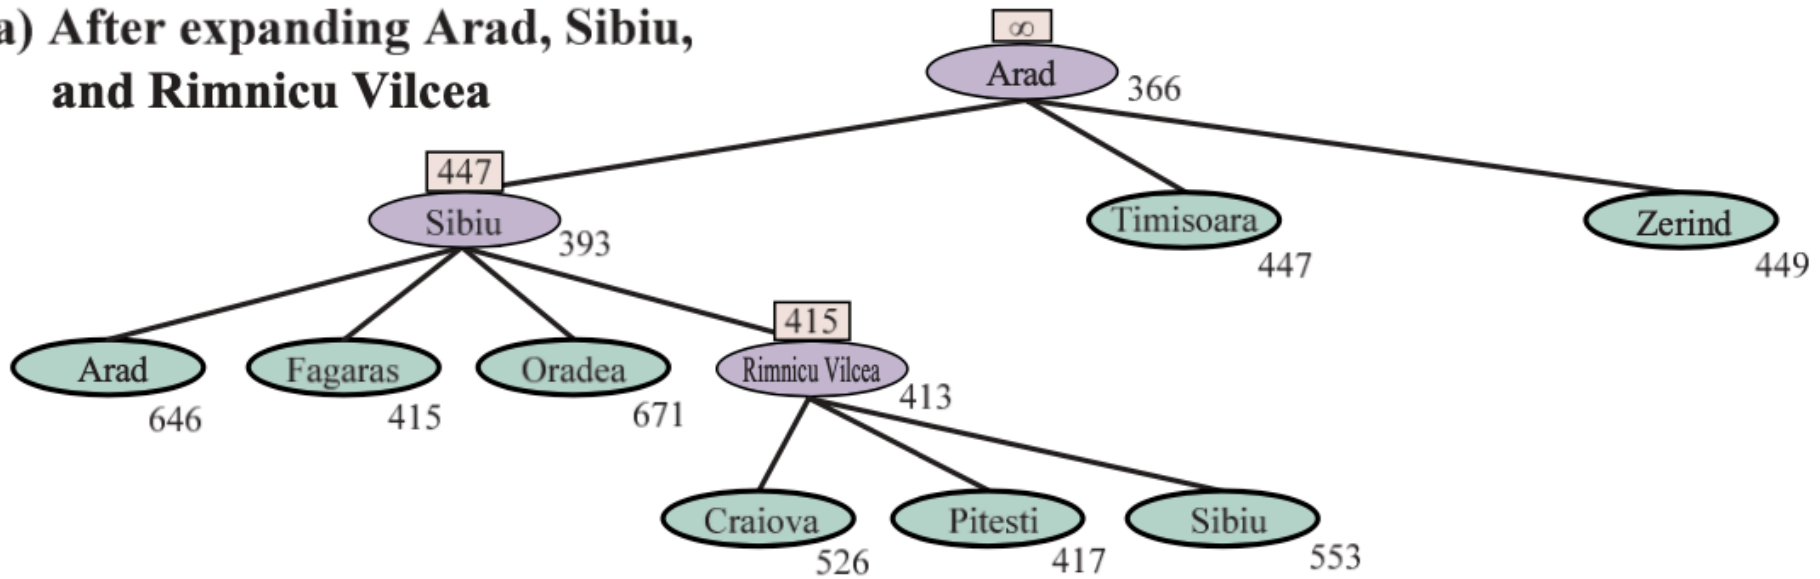
Algorithm

function RECURSIVE-BEST-FIRST-SEARCH(*problem*) **returns** a solution or *failure*
 solution, fvalue \leftarrow RBFS(*problem*, NODE(*problem*.INITIAL), ∞)
return *solution*

function RBFS(*problem, node, f_limit*) **returns** a solution or *failure*, and a new *f*-cost limit
 if *problem*.IS-GOAL(*node*.STATE) **then return** *node*
 successors \leftarrow LIST(EXPAND(*node*))
 if *successors* is empty **then return** *failure, ∞*
 for each *s* **in** *successors* **do** // update *f* with value from previous search
 s.f \leftarrow max(*s*.PATH-COST + *h*(*s*), *node.f*)
 while true do
 best \leftarrow the node in *successors* with lowest *f*-value
 if *best.f* > *f_limit* **then return** *failure, best.f*
 alternative \leftarrow the second-lowest *f*-value among *successors*
 result, best.f \leftarrow RBFS(*problem, best, min(f_limit, alternative)*)
 if *result* \neq *failure* **then return** *result, best.f*

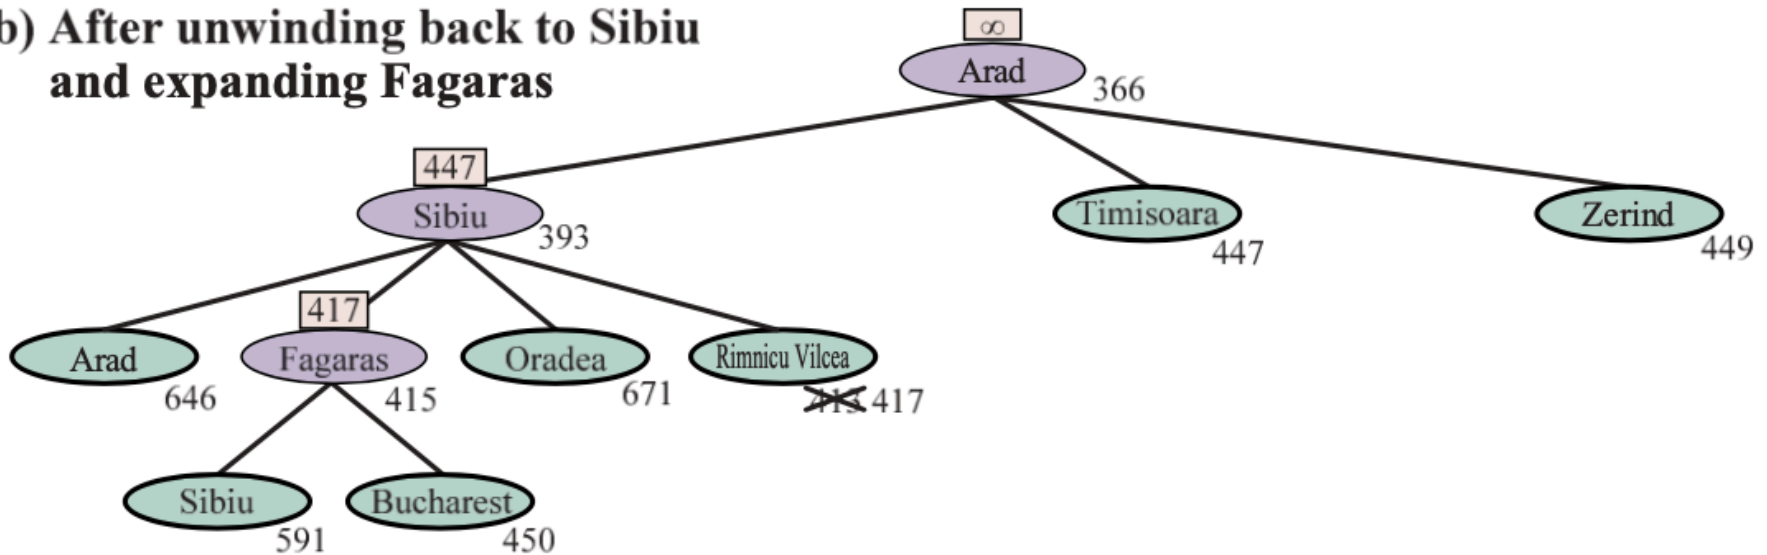
Recursive Best-First Search (RBFS)

(a) After expanding Arad, Sibiu, and Rimnicu Vilcea



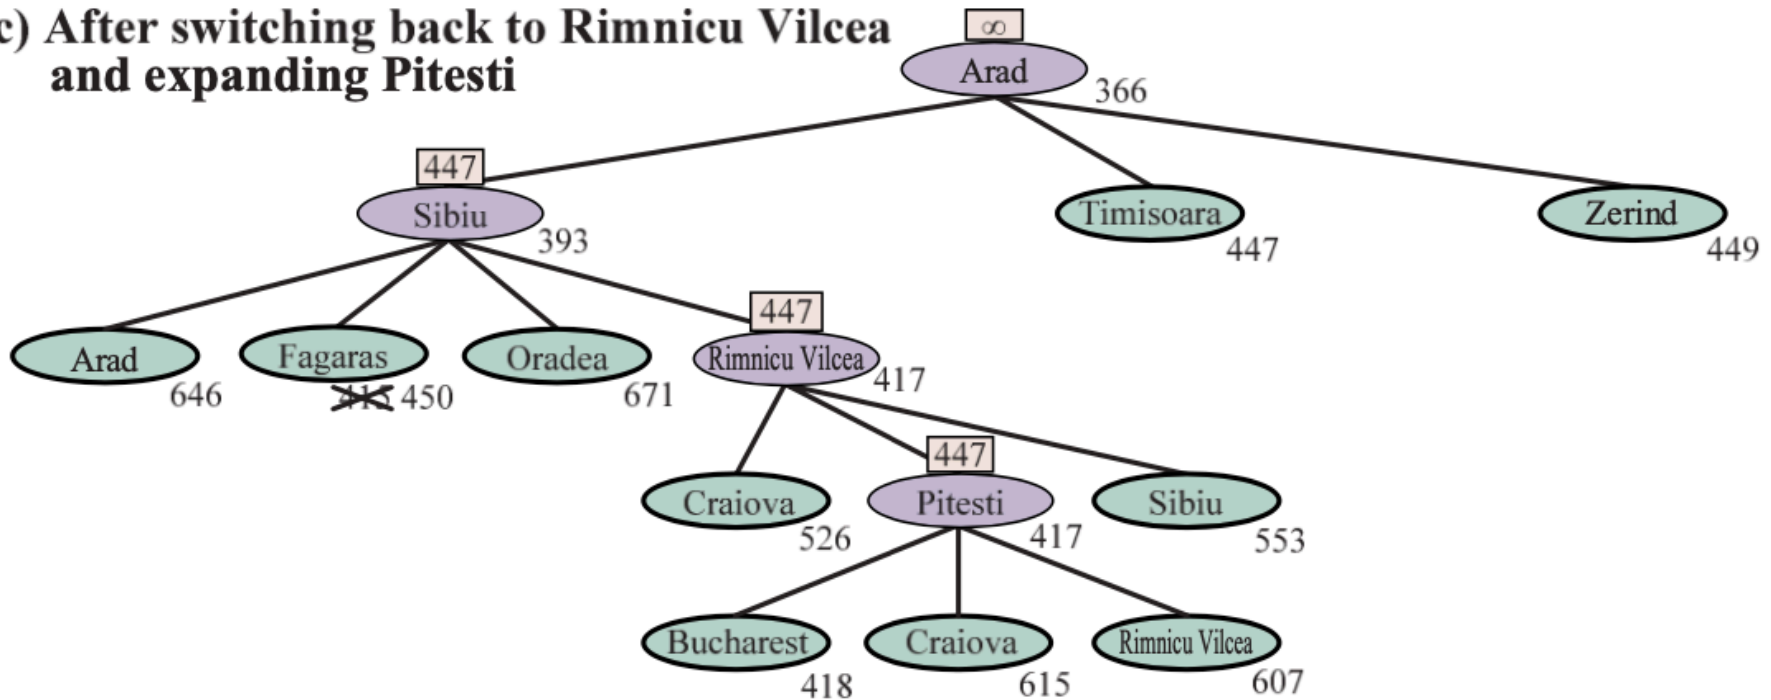
Recursive Best-First Search (RBFS)

(b) After unwinding back to Sibiu and expanding Fagaras

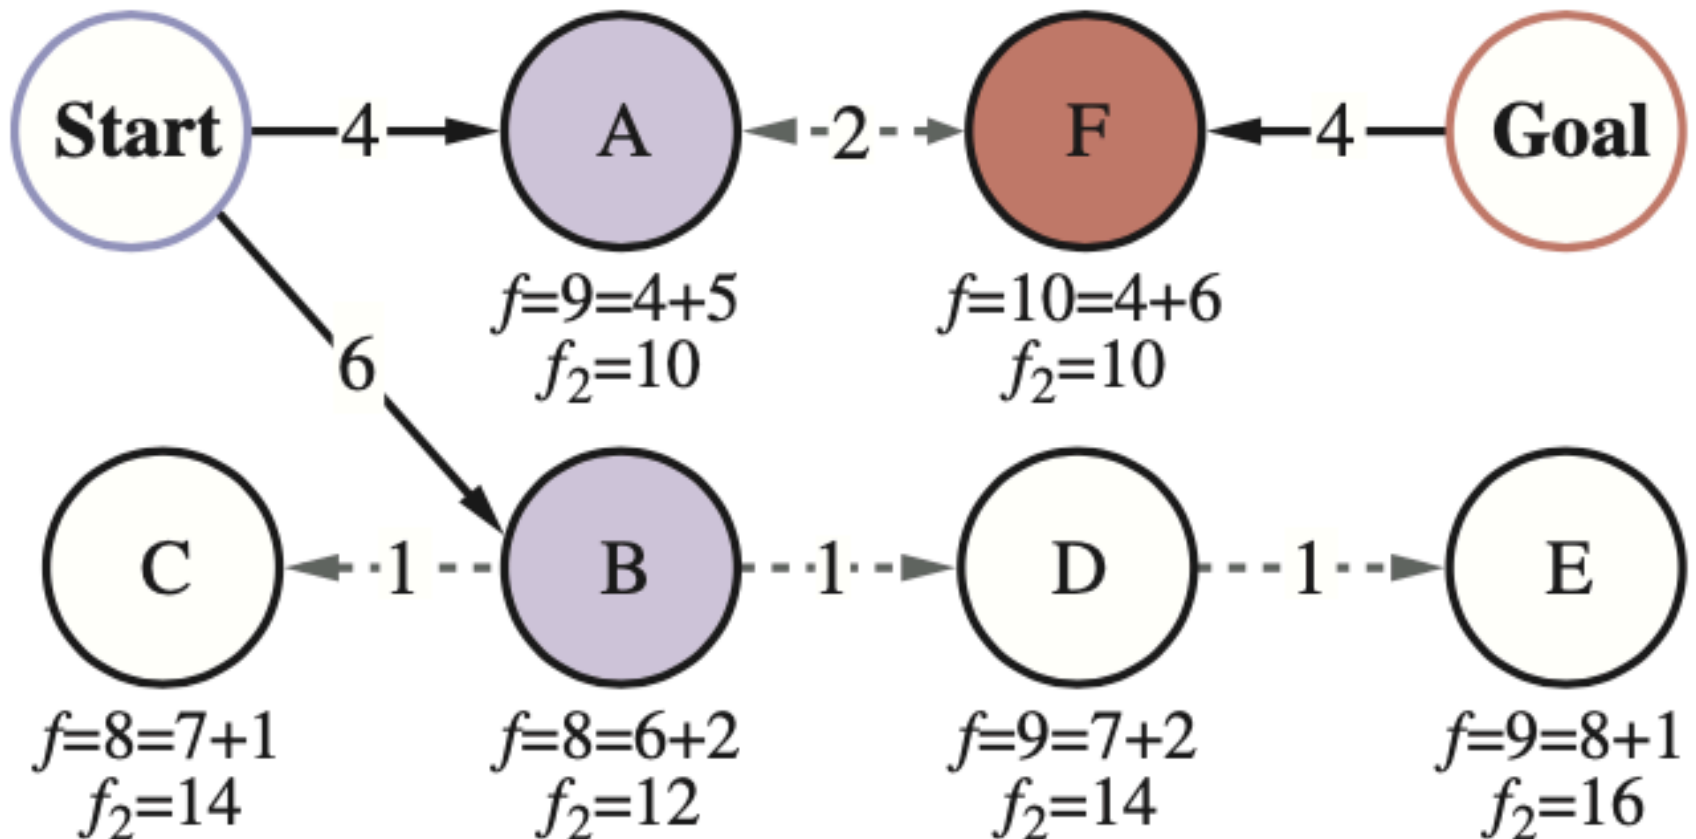


Recursive Best-First Search (RBFS)

(c) After switching back to Rimnicu Vilcea and expanding Pitesti



Bidirectional Search maintains two frontiers



On the left, nodes A and B are successors of Start;
on the right, node F is an inverse successor of Goal

A typical instance of the 8-puzzle

The shortest solution is 26 actions long

| | | |
|---|---|---|
| 7 | 2 | 4 |
| 5 | | 6 |
| 8 | 3 | 1 |

Start State

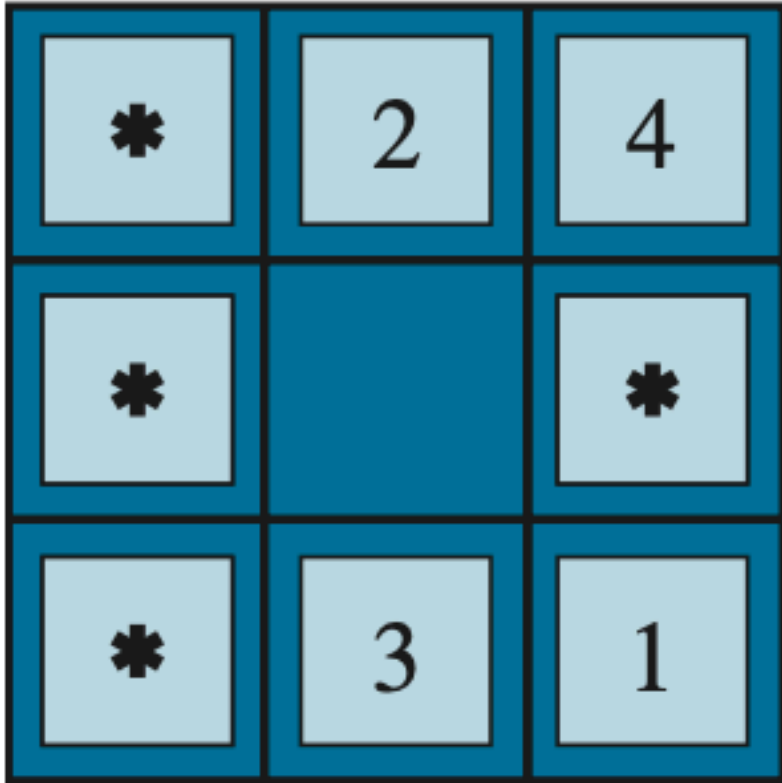
| | | |
|---|---|---|
| | 1 | 2 |
| 3 | 4 | 5 |
| 6 | 7 | 8 |

Goal State

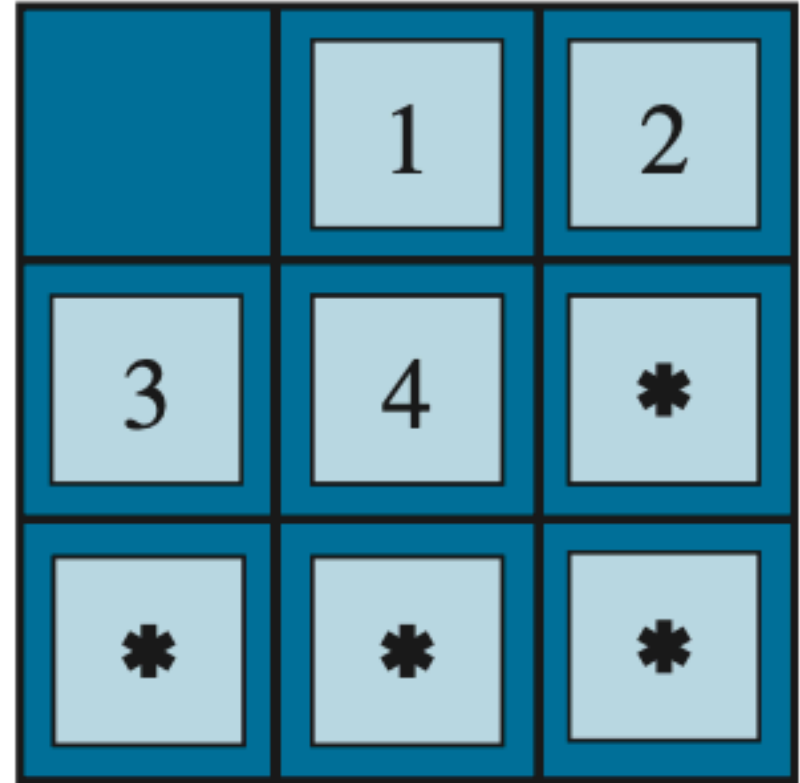
Comparison of the search costs and effective branching factors for 8-puzzle problems

| d | Search Cost (nodes generated) | | | Effective Branching Factor | | |
|-----|-------------------------------|------------|------------|----------------------------|------------|------------|
| | BFS | $A^*(h_1)$ | $A^*(h_2)$ | BFS | $A^*(h_1)$ | $A^*(h_2)$ |
| 6 | 128 | 24 | 19 | 2.01 | 1.42 | 1.34 |
| 8 | 368 | 48 | 31 | 1.91 | 1.40 | 1.30 |
| 10 | 1033 | 116 | 48 | 1.85 | 1.43 | 1.27 |
| 12 | 2672 | 279 | 84 | 1.80 | 1.45 | 1.28 |
| 14 | 6783 | 678 | 174 | 1.77 | 1.47 | 1.31 |
| 16 | 17270 | 1683 | 364 | 1.74 | 1.48 | 1.32 |
| 18 | 41558 | 4102 | 751 | 1.72 | 1.49 | 1.34 |
| 20 | 91493 | 9905 | 1318 | 1.69 | 1.50 | 1.34 |
| 22 | 175921 | 22955 | 2548 | 1.66 | 1.50 | 1.34 |
| 24 | 290082 | 53039 | 5733 | 1.62 | 1.50 | 1.36 |
| 26 | 395355 | 110372 | 10080 | 1.58 | 1.50 | 1.35 |
| 28 | 463234 | 202565 | 22055 | 1.53 | 1.49 | 1.36 |

A subproblem of the 8-puzzle



Start State



Goal State

The task is to get tiles 1, 2, 3, 4, and the blank into their correct positions, without worrying about what happens to the other tiles

A Web service providing driving directions, computed by a search algorithm.

7 h 17 min (578 km)

via A1 and E81

Fastest route, the usual traffic

Arad

Romania

> Get on A1 from Strada Andrei Șaguna, Strada Dorobanți and Calea Bodroglu/DJ682F

8 min (4.3 km)

> Get on DN68A/E673 in Județul Timiș from A1

53 min (101 km)

> Get on A1 in Șoimăș from DN68A/E673 and DN7/E68

1 h 13 min (75.7 km)

> Follow A1, DN1/DN7/E68/E81 and A1 to DN1/DN7/E68/E81 in Județul Sibiu. Take the DN7/DN1 exit from A1

1 h 18 min (131 km)

> Continue to Râmnicu Vâlcea

1 h 21 min (95.4 km)

↑ Continue onto DN7/E81

Continue to follow E81

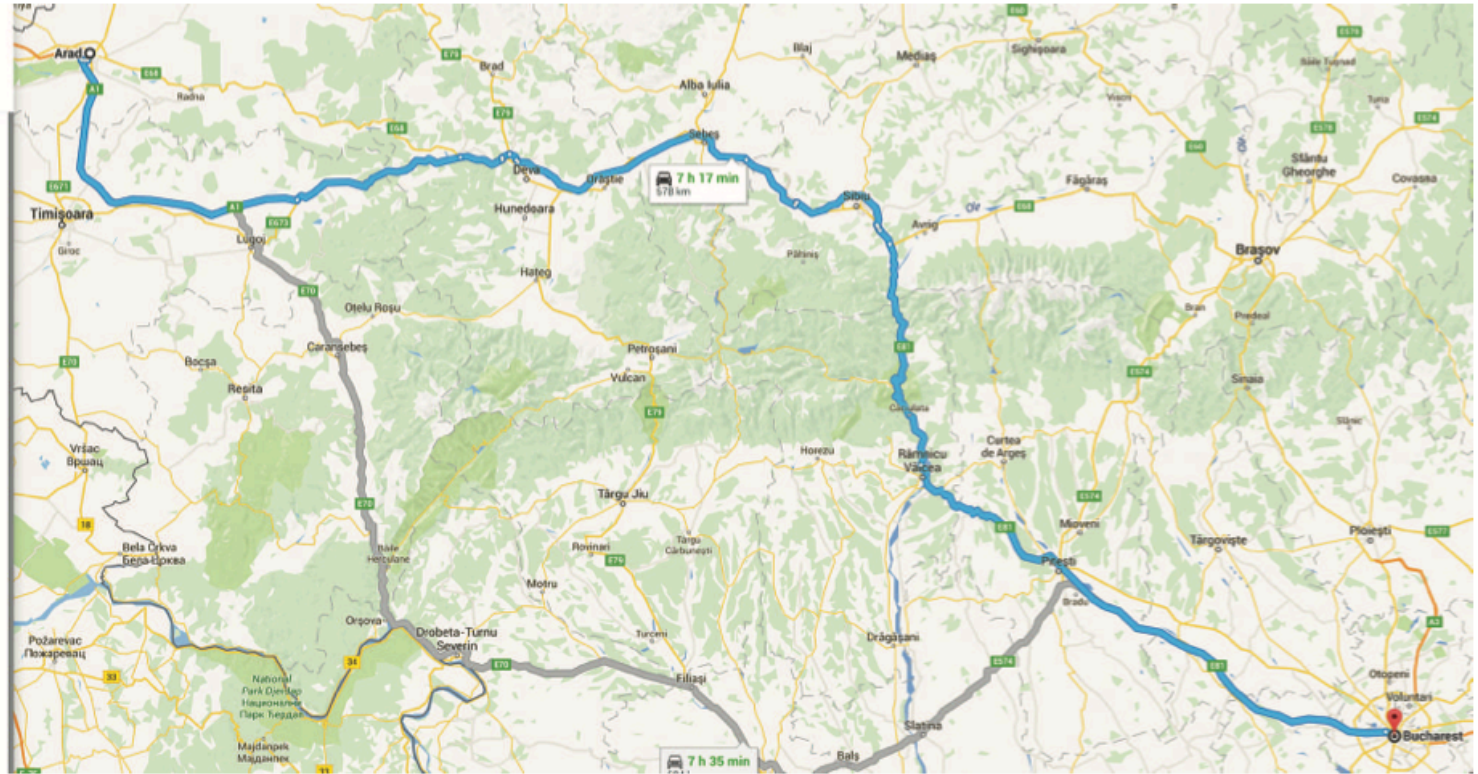
2 h 15 min (177 km)

> Take Siplau Independenței to Bulevardul Unirii/E81

9 min (3.5 km)

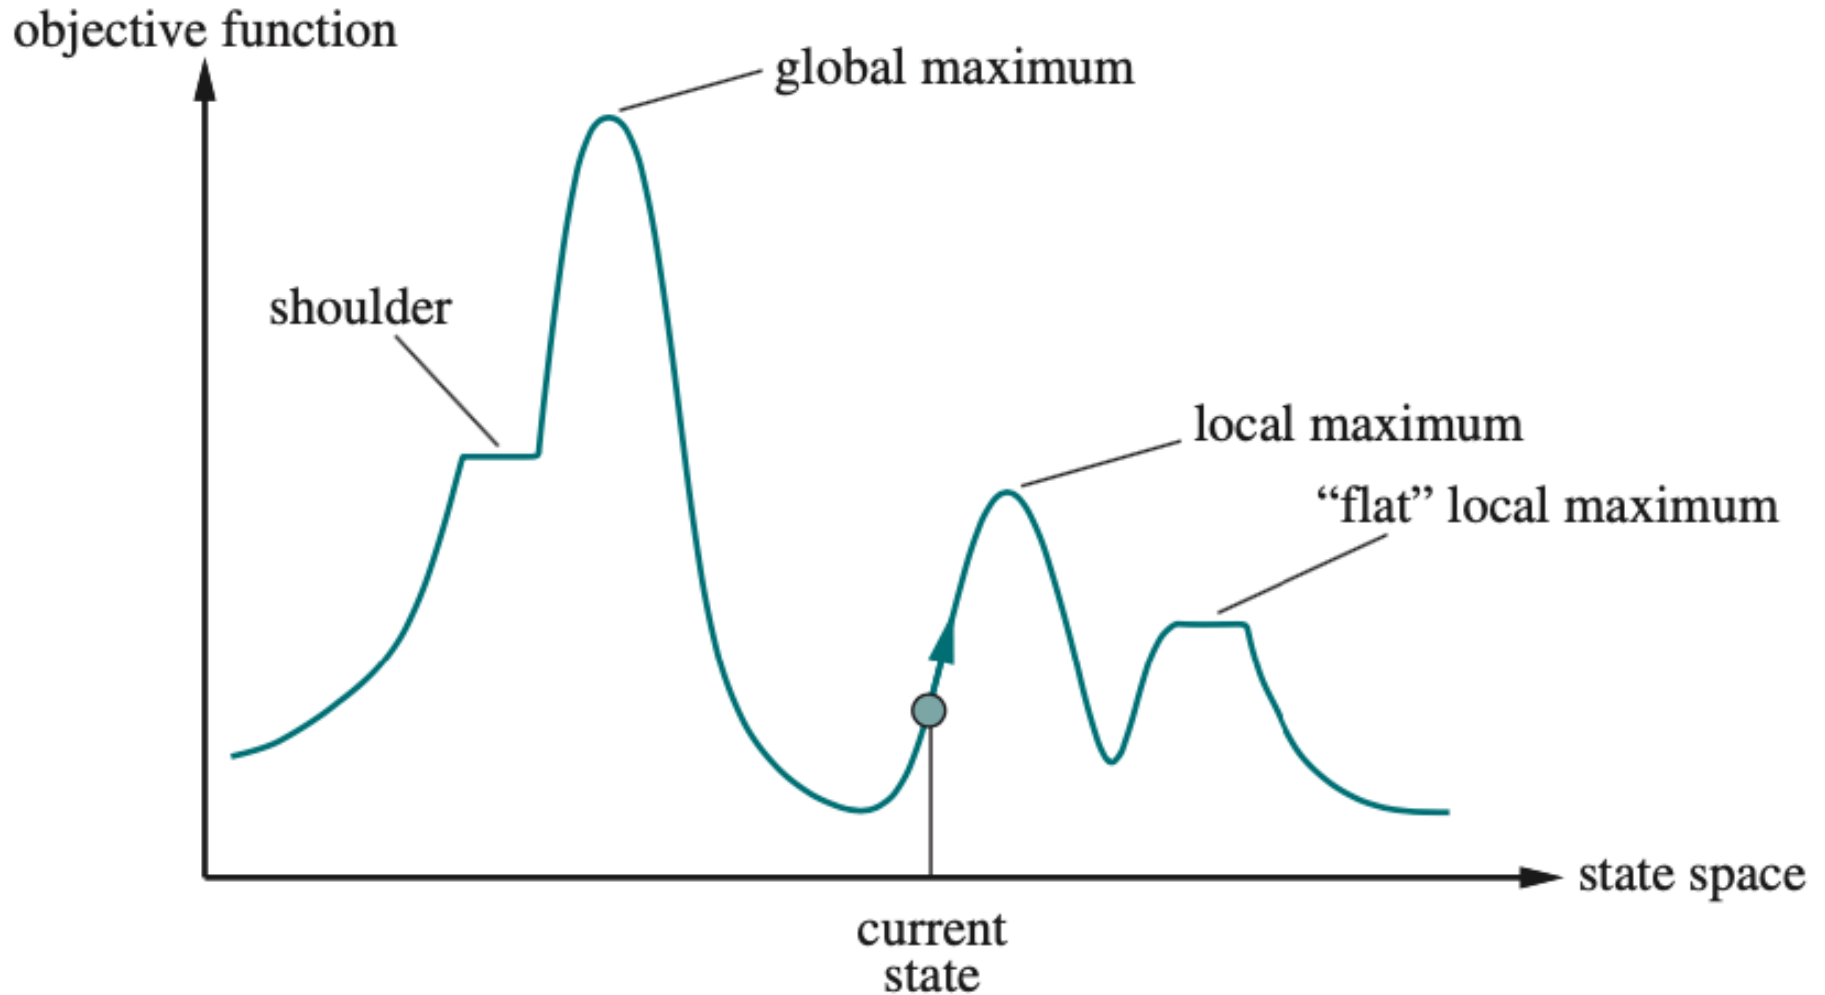
Bucharest

Romania



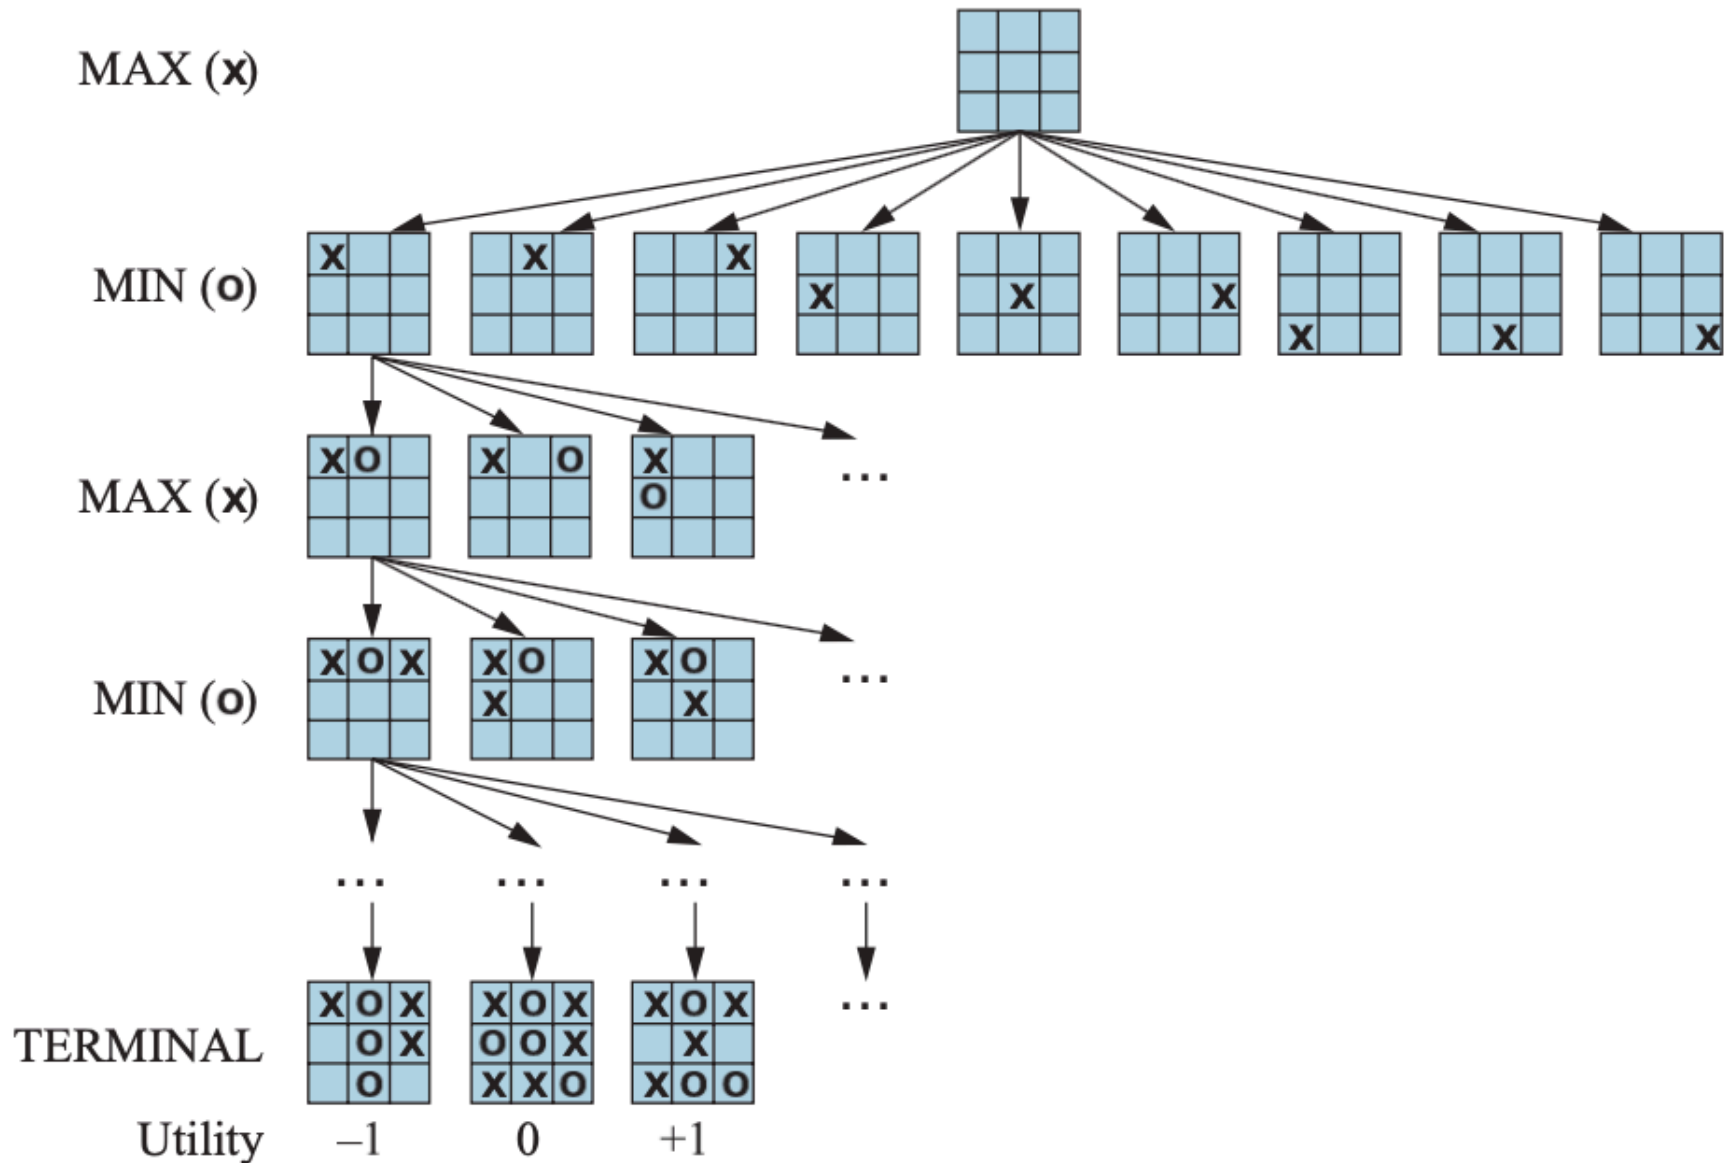
Search in Complex Environments

A one-dimensional state-space landscape



Adversarial Search and Games

Game Tree for the Game of Tic-tac-toe

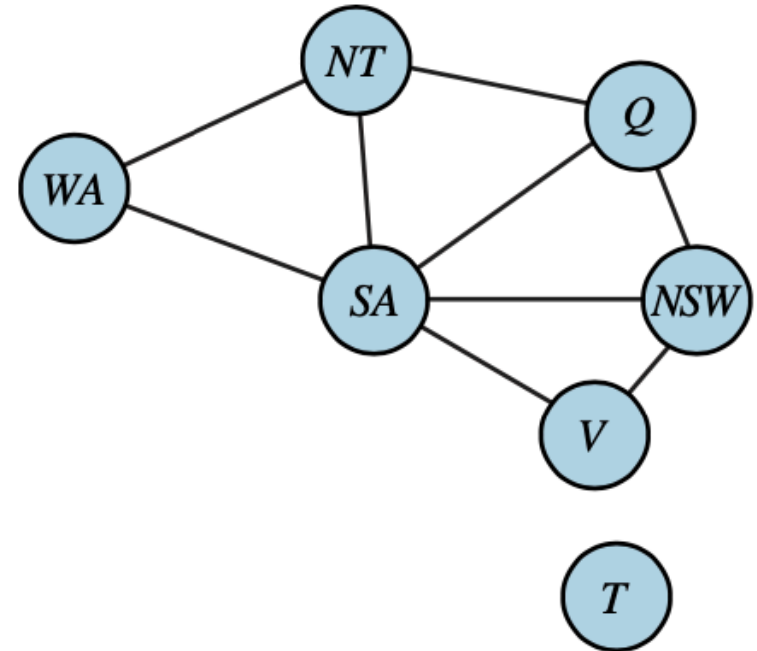


Constraint Satisfaction Problems

The Map-Coloring Problem Represented as a Constraint Graph

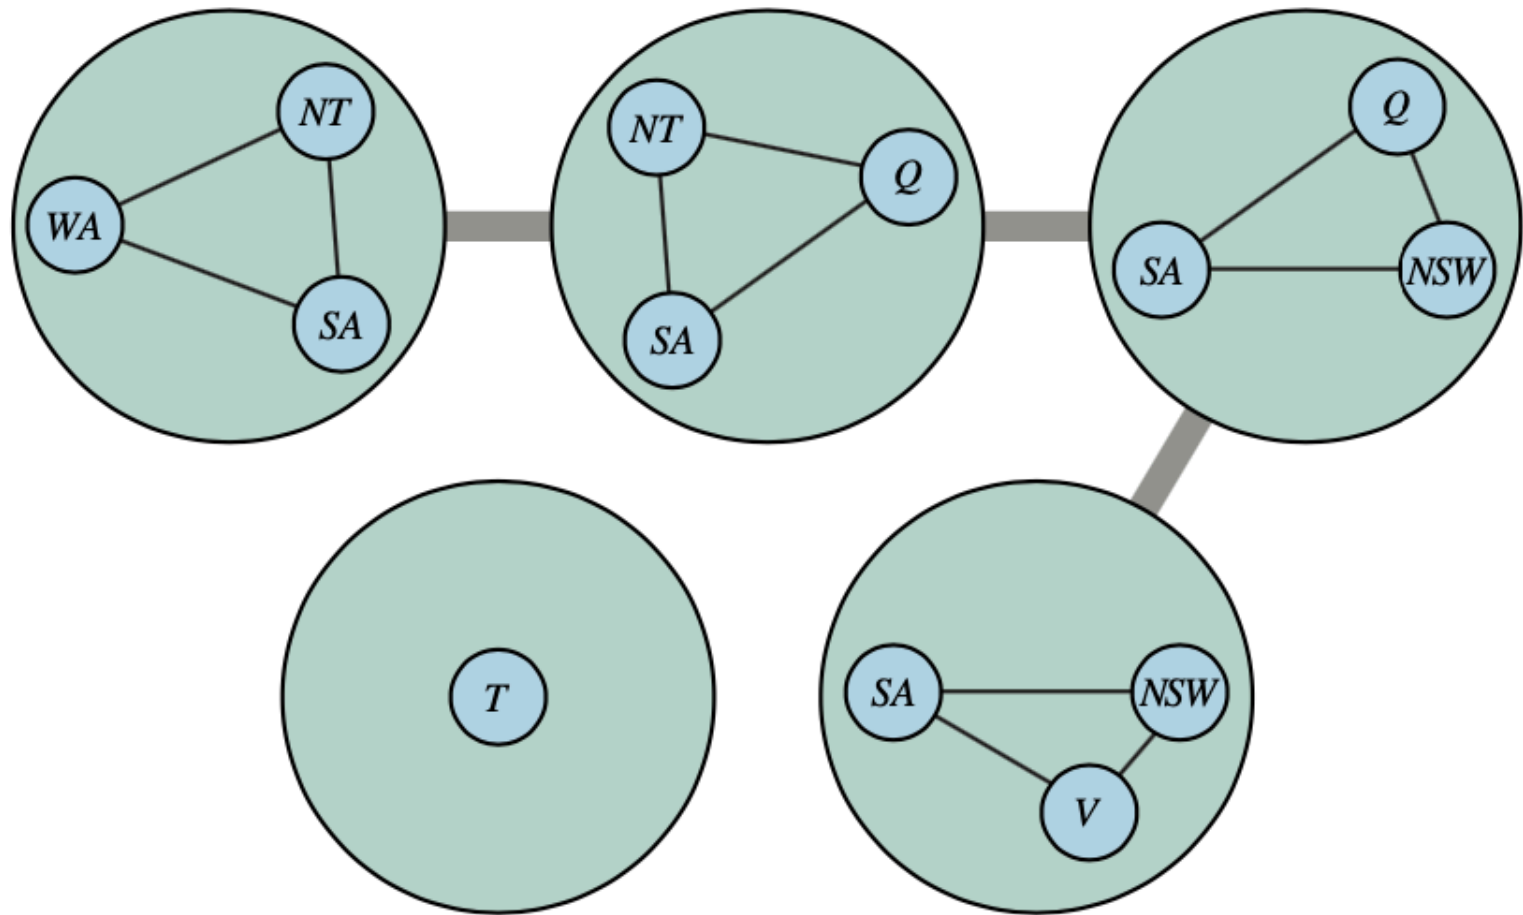


(a)



(b)

A Tree Decomposition of the Constraint Graph

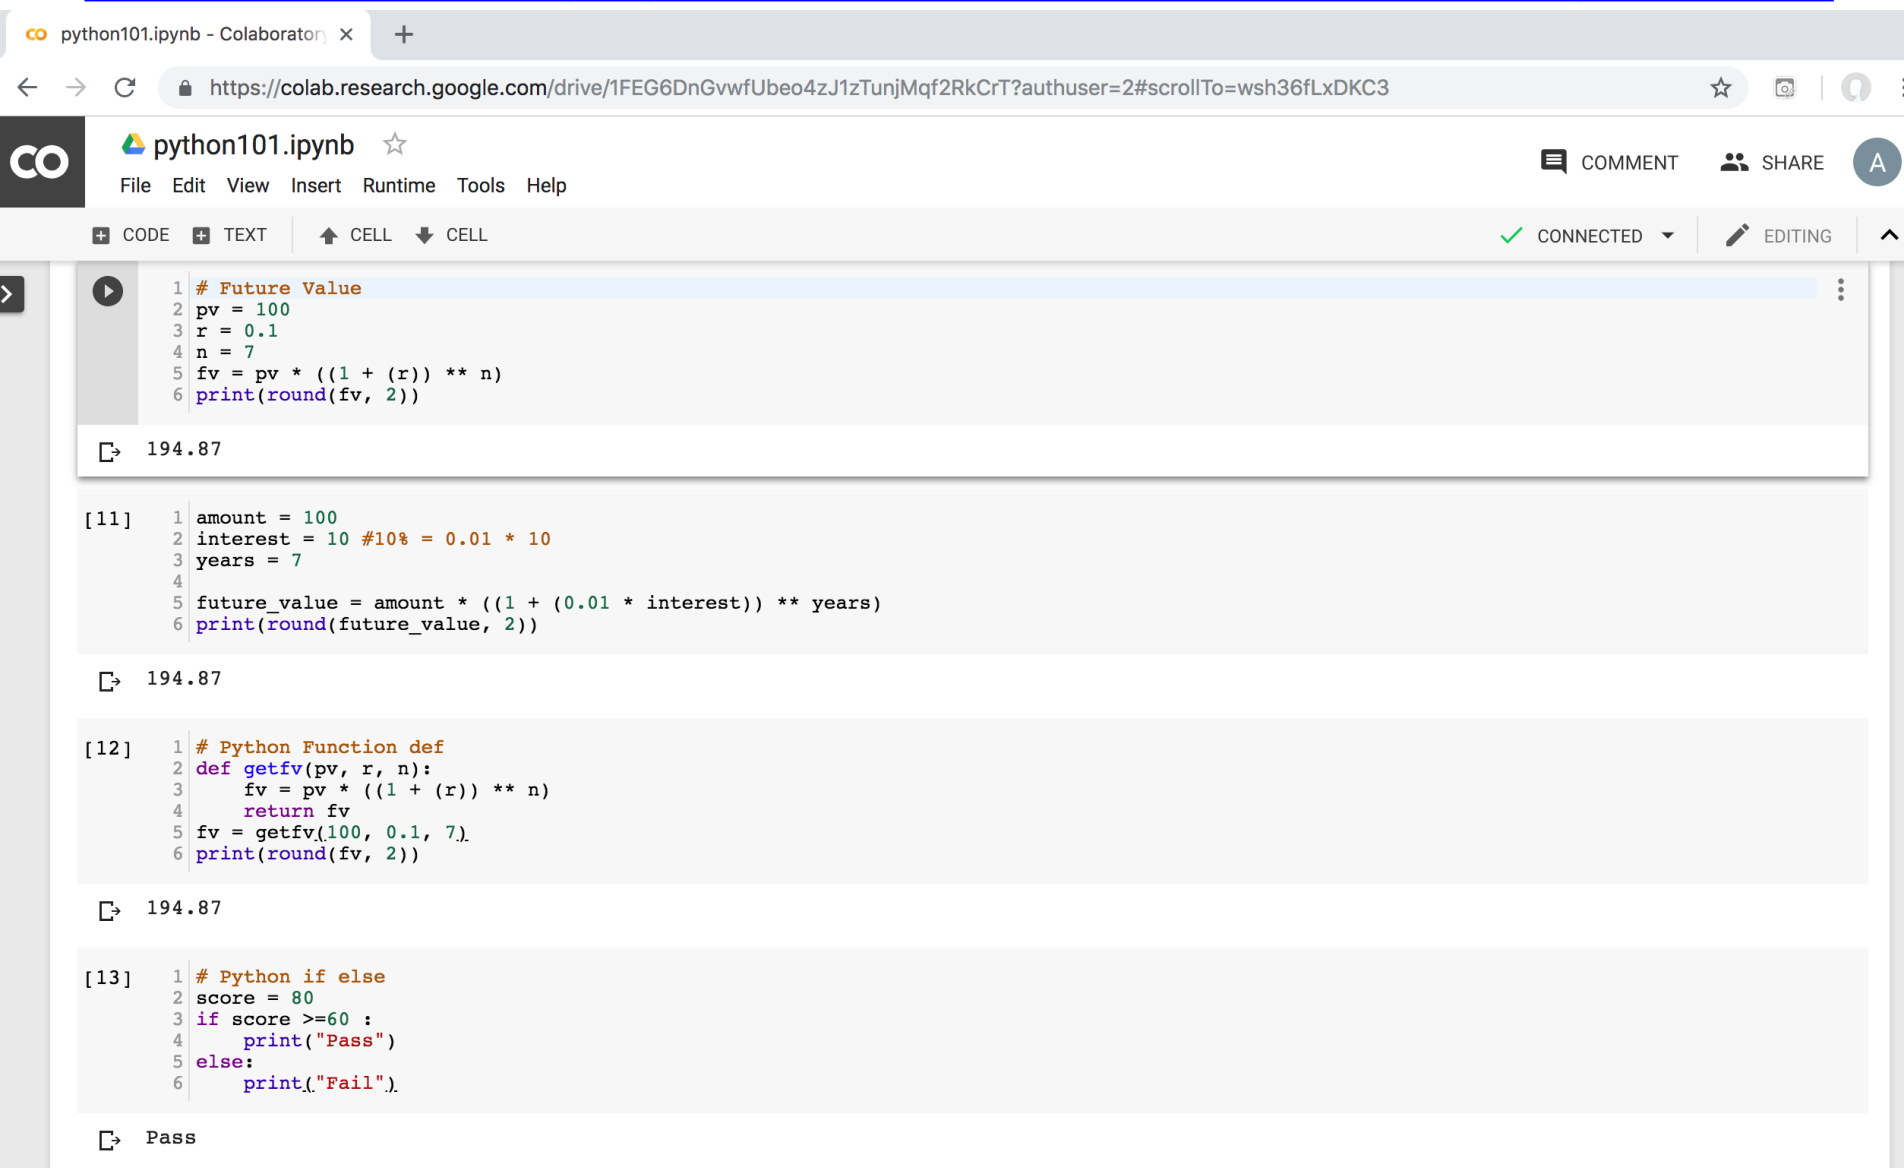


AIMA Python

- Artificial Intelligence: A Modern Approach (AIMA)
 - <http://aima.cs.berkeley.edu/>
- AIMA Python
 - <http://aima.cs.berkeley.edu/python/readme.html>
- Search
 - <http://aima.cs.berkeley.edu/python/search.html>
- Games: Adversarial Search
 - <http://aima.cs.berkeley.edu/python/games.html>
- CSP (Constraint Satisfaction Problems)
 - <http://aima.cs.berkeley.edu/python/csp.html>

Python in Google Colab (Python101)

<https://colab.research.google.com/drive/1FEG6DnGvwfUbeo4zJ1zTunjMqf2RkCrT>



python101.ipynb - Colaboratory

File Edit View Insert Runtime Tools Help

COMMENT SHARE

CONNECTED EDITING

```
1 # Future Value
2 pv = 100
3 r = 0.1
4 n = 7
5 fv = pv * ((1 + (r)) ** n)
6 print(round(fv, 2))
```

194.87

```
[11] 1 amount = 100
2 interest = 10 #10% = 0.01 * 10
3 years = 7
4
5 future_value = amount * ((1 + (0.01 * interest)) ** years)
6 print(round(future_value, 2))
```

194.87

```
[12] 1 # Python Function def
2 def getfv(pv, r, n):
3     fv = pv * ((1 + (r)) ** n)
4     return fv
5 fv = getfv(100, 0.1, 7).
6 print(round(fv, 2))
```

194.87

```
[13] 1 # Python if else
2 score = 80
3 if score >=60 :
4     print("Pass")
5 else:
6     print("Fail").
```

Pass

<https://tinyurl.com/aintpupython101>

Summary

- **Solving Problems by Searching**
- **Search in Complex Environments**
- **Adversarial Search and Games**
- **Constraint Satisfaction Problems**

References

- Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson.
- Aurélien Géron (2019), Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems, 2nd Edition, O'Reilly Media.